AFIT/GE/ENG/99M-30

GPS Signal Offset Detection and Noise Strength Estimation in a Parallel Kalman Filter Algorithm

THESIS

Barry J. Vanek
Second Lieutenant, USAF

AFIT/GE/ENG/99M-30

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. AGENCY USE ONLY *(Leave blank)*	2. REPORT DATE	3. REPORT TYPE AND DATES COVERED
	March 1999	MS Thesis

4. TITLE AND SUBTITLE
GPS Signal Offset Detection and Noise Strength Estimation in a Parallel Kalman Filter Algorithm

5. FUNDING NUMBERS

6. AUTHOR(S)
Barry J. Vanek

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
Air Force Institute Technology, WPAFB OH 45333-6583

8. PERFORMING ORGANIZATION REPORT NUMBER
AFIT/GE/ENG/99M-30

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)
Juan Vasquez, Capt, USAF
AFRL/SNAR Bldg 620
2241 Avionics Circle
WPAFB, OH 45433-7321
(937) 255-5668 x4014

10. SPONSORING/MONITORING AGENCY REPORT NUMBER

11. SUPPLEMENTARY NOTES

12a. DISTRIBUTION AVAILABILITY STATEMENT
Distribution Unlimited

12b. DISTRIBUTION CODE

13. ABSTRACT *(Maximum 200 words)*

Measurements from Global Positioning System (GPS) satellites are subject to corruption by signal interference and induced offsets. This thesis presents two independent algorithms to ensure the navigation system remains uncorrupted by these possible GPS failures. The first is a parameter estimation algorithm that estimates the measurement noise variance of each satellite. A redundant measurement differencing (RMD) technique provides direct observability of the differenced white measurement noise samples. The variance of the noise process is estimated and provided to the second algorithm, a parallel Kalman filter structure, which then adapts to changes in the real-world measurement noise strength.

The parallel Kalman filter structure detects and isolates signal offsets in individual GPS satellites. The offset detection algorithm calculates test statistics on each of the filters and makes decisions on whether to remove satellites from the solution based on these statistics. The two algorithms contain several user-defined parameters that have significant effects when adjusted. The various effects of parameter variation are described and a parameter set is chosen at which to evaluate the algorithms. The combined algorithm performs quite well in computer simulations.

14. SUBJECT TERMS	15. NUMBER OF PAGES
Global Positioning System; Kalman Filtering; Error Detection, Isolation, and Estimation GPS; INS; Adaptive Filtering; Chi-square	98
	16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT
Unclassified	Unclassified	Unclassified	UL

Disclaimer

The views expressed in this thesis are those of the author and do not reflect the official policy or position of the United States Air Force, the Department of Defense, or the United States Government.

AFIT/GE/ENG/99M-30

GPS Signal Offset Detection and Noise Strength Estimation in a Parallel Kalman Filter Algorithm

THESIS

Presented to the Faculty of the Graduate School of Engineering of the Air Force Institute of Technology Air University In Partial Fulfillment for the Degree of

Master of Science in Electrical Engineering

Barry J. Vanek

Second Lieutenant, USAF

Air Force Institute of Technology

Wright-Patterson AFB, Ohio

March 8, 1999

Sponsored in part by the Sensors Directorate, Air Force Research Laboratory

AFIT/GE/ENG/99M-30

GPS Signal Offset Detection and Noise Strength Estimation in a Parallel Kalman Filter Algorithm

Barry J. Vanek

Second Lieutenant, USAF

Approved:

Peter S. Maybeck — 11 Mar 99

Peter S. Maybeck, Committee Chairman — Date
Professor of Electrical Engineering
Department of Electrical and Computer Engineering

John F. Raquet — 12 MAR 99

John F. Raquet, Committee Member — Date
Captain, USAF
Assistant Professor of Electrical Engineering
Department of Electrical and Computer Engineering

Dedication

I dedicate this work to my family, for always listening to my technobabble without yawning, hanging up, or changing their telephone numbers. Mom and her constant support. Dad and his endless belief in his sons. Steve for always being there when we were growing up. And Mike for reminding me to keep my reality and my dreams one in the same.

Acknowledgments

I would like to start by acknowledging an old high school teacher of mine who started me on the road to engineering and science. I don't think of him as a mentor, yet I can definitely trace my interest in the sciences back to him. I don't thank him for what he taught me, but for everything he made me figure out. I can't remember him ever answering a question of mine, not one. He did ask a lot of questions, which made me want answers. So, here I am starting to answer a couple of questions. Thank you, Mr. Vondruska.

I want to recognize the Air Force for sending me through this program. Without AFIT, it would have been a long time before I attempted a full-time graduate program in engineering. I guess I hit paydirt in the AFPC dice game. Also, somewhat after fact, the Air Force deserves thanks for the great education I received at Iowa State University. Without that scholarship, I would not have experienced five years of stress, fun, and life.

I want to thank the professors here at AFIT. I'm not sure, exactly, what is attractive about a DOD operated education institute, but whatever that is for each of you, I'm glad you are here. I especially want to acknowledge my thesis advisor, Dr. Peter Maybeck. He kept me looking forward and up through some arduous times. A special thanks goes out to Stan Musick who selflessly assists students in need and to Don Smith who has the never ending task of maintaining the computer systems in our lab.

A large amount of appreciation goes out to all my friends at AFIT. This isn't said lightly, because I tend to be brash, opinionated, and somewhat annoying at times. I owe most of them an apology for all those questions I asked right before break. The time spent here has been more than just work. In that spirit, my thanks goes out to W.O. Wright's and Waffle House for staying open while the rest of Dayton is snoring.

Table Of Contents

List of Figures

List of Tables

Abstract

Measurements from Global Positioning System (GPS) satellites are subject to corruption by signal interference and induced offsets. This thesis presents two independent algorithms to ensure the navigation system remains uncorrupted by these possible GPS failures. The first is a parameter estimation algorithm that estimates the measurement noise variance of each satellite. A redundant measurement differencing (RMD) technique provides direct observability of the differenced white measurement noise samples. The variance of the noise process is estimated and provided to the second algorithm, a parallel Kalman filter structure, which then adapts to changes in the real-world measurement noise strength.

The parallel Kalman filter structure detects and isolates signal offsets in individual GPS satellites. The offset detection algorithm calculates test statistics on each of the filters and makes decisions on whether to remove satellites from the solution based on these statistics. The two algorithms contain several user-defined parameters that have significant effects when adjusted. The various effects of parameter variation are described and a parameter set is chosen at which to evaluate the algorithms. The combined algorithm performs quite well in computer simulations.

GPS Signal Offset Detection and Noise Strength Estimation in a Parallel Kalman Filter Algorithm

Chapter 1 - Introduction

1.1 Background

Navigation systems are used on most aircraft, both military and commercial. Autonomous landing systems, tactical delivery of precision weapons, and reference systems at test facilities all require a high degree of positioning accuracy. In high precision navigation systems, an inertial navigation system (INS) is aided by a Kalman filter (KF) that processes external measurements. Global Positioning System (GPS) is the most widely used, high accuracy sensor that provides the KF with measurements. Because GPS is partially space-based, the system parameters and error characteristics are functions of many variables and can vary greatly between uses. Generally, GPS receivers use very simple models, if any at all, for compensation of errors.

Today's navigation applications demand high confidence in their solution. Often, many lives depend on it. When GPS is corrupted by intentional or naturally induced errors, the KF in most systems cannot adapt to, or even detect, these changes. Sometimes it becomes better not to use the corrupted measurements at all. The confidence placed in a precise navigation solution is directly related to the integrity of the data used in that solution. The estimation and detection algorithms described in this thesis could potentially protect navigation systems from corruption by intentional and/or natural errors in GPS signals. With this protection, confidence increases and bold innovations can move forward. Autonomous flight vehicles could rely more heavily on the navigation solution during precision maneuvers, such as landing on an aircraft carrier or flying in tight formation. More importantly, the aging Instrument Landing System is due to be replaced by a GPS-based Precision Landing System (PLS) [6]. Even if the accuracy constraints can be met with some aided-GPS system, the need for confidence in that accurate solution is extremely high.

1.2 Problem Definition

This thesis focuses on two main ideas: (1) estimating the noise variance on incoming GPS measurements and adaptively tuning the KF in real-time; and (2) developing a parallel KF structure which detects bias-like errors in GPS measurements, isolates the time of error initiation, and resets the main KF to an uncorrupted solution. During this development, the goal is always real-world implementation, using assumptions that are as realistic as possible. Unreasonable demands that would require significant modifications to hardware or aircraft are not made. Extra effort is taken to make the application of this work to modern navigation systems as smooth and quick as possible. A majority of the modifications can be made to the navigation software without any hardware requirements.

1.3 Scope

This work is strictly a proof of concept. Representative low order integration filters are used in this study. Performance evaluations are shown, but only for comparison to a baseline, non-adaptive system. A single profile of steady-level flight is examined. Two types of measurement errors, signal offsets (biases or ramps) and increased noise variances, are induced at varying magnitudes and magnitude rates. A bias can represent a spoof[1], multipath, ephemeris error, or perhaps an unmodeled or uncompensated atmospheric delay. This last type of error is unlikely to occur in a high precision receiver, although the solar maximum in the coming years could create significant deviations from standard half-cosine models [15]. Each GPS satellite is treated separately for both signal offset detection and noise variance estimation. By handling each pseudorange (PR) separately, bad satellites can be indicated and removed from the solution independently of other satellites. Increased measurement noise can represent in-band or out-of-band interference, intentional jamming, degraded satellites, or satellites low on the horizon.

Code-only range measurements are input to the KF. This research will not simulate phase measurements or carrier phase receivers, since they are not yet standard equipment on military or

[1]Intentional biases induced by outside sources are referred to as spoofs.

civilian aircraft. Range-rate measurements are commonly output by GPS receivers, but are not used by in this simulation. An advanced spoof signal would adjust phase [7] as well as code, but for a first cut at error detection, only code is examined. This study should apply, without modification, to the new Block IIF GPS signal configuration since the algorithm only requires range measurements. Stated generally, the algorithms shown in this work are independent of the measurement device under scrutiny. The device only needs to be a filter receiving measurements which help it estimate system errors that grow over time. This independent nature stems from the level on which the algorithms perform, i.e. measurement incorporation rather than signal processing.

Considerable effort is given to maintaining applicability to commercial aviation while modeling a military system. Reasonable component specifications are selected to apply to either military or commercial aircraft. The INS modelled in the study has one standard deviation (1σ) horizontal error growth of 0.42 nmi/hr, representative of many current INSs. The military ability to receive dual-frequency P-code measurements without selective availability (SA) separates it from commercial aviation. Differential corrections, when applied to C/A-code measurements, eliminate the effects of SA. Standard differential GPS requires a ground station. It is reasonable to assume any airport, military or commercial, has or soon will have access to differential corrections. For military operations in unfriendly territory, it is common for ground units to set up temporary differential stations. A measurement noise level is chosen to be representative of a dual-frequency P-code receiver receiving differential corrections.

Four satellites are considered in-view at all time. Typical GPS receivers require four satellites to provide sufficient observability to estimate a three-dimensional position and the user clock error. Including more satellites would only increase the navigation accuracy and the time it takes to test the incoming measurements for errors. When an offset is detected in one satellite, only three satellites are included in the solution. In a real-world application, it would be desirable to test the measurements from as many satellites as possible and when a failure is declared, a new set of satellites are selected to use in the solution.

The simulations in this research take place in the presence of an integrated navigation system composed of an INS with a barometric altimeter, and a GPS receiver. The barometric altimeter is included primarily to bound the INS's unstable vertical channel. Other aids such as pseudolites or radar altimeters are excluded from this study, as it is only a proof of concept. These additional aids could be added at a later time to evaluate the algorithm's performance in a high accuracy navigation system. It should be noted that the theoretical limit of error detection is determined by the accuracy of the combined navigation system without GPS, i.e. the better the INS and aids other than GPS, the smaller the detectable error in the GPS measurements. This idea is developed further in Chapter 3.

1.4 System Description and Assumptions

An extended Kalman filter is used to estimate errors in the INS by incorporating measurements from aiding sources, viz. GPS. There are significant benefits to using error states instead of total states. The INS errors are better modeled with a linear perturbation model than are aircraft dynamics [2]. KF performance is less sensitive to varying dynamics when formulated in error states than in total states. INS errors change slowly with time; therefore, a longer filter propagation time (integration sub-interval size) can be used without creating significant inaccuracies with an error state filter. A short propagation time is needed in a total state formulation when the values of the variables significantly move away from the point about which the linear approximation is made.

The GPS/INS integration is implemented in a tightly coupled, error state formulation. Tight coupling of a GPS receiver to an INS is the application of individual, raw pseudoranges as separate measurements to the KF. Loose coupling is when GPS's navigation solution, calculated by its own KF, is applied as a single measurement to a KF modeling the INS errors. For GPS to calculate a three-dimensional position solution independently, four satellites are required. Fewer than four forces the receiver to make assumptions which can significantly degrade the precision of the solution. In a tight configuration, any number of satellites can be used and their measurements are properly applied in the filter. This allows for the isolation of an error to a specific satellite and its elimination

from use in the solution. After the bad satellite is detected and its signal removed from use in the receiver, another satellite's measurements, known to be good, can be brought into the filter. It would be possible to monitor all the satellites in view since most modern GPS receivers have twelve channels to track GPS satellites.

The work in this thesis uses a feedforward implementation of the KF estimates to the INS solution. This configuration simply subtracts the KF error estimate in a calculation from the INS output. This method prevents measurements from possibly corrupting the INS itself. In a feedback configuration the INS platform, either a true platform or just a reference in a computer, is reset after every update. The FAA currently restricts feedback configurations due to the possibility of corrupting the INS output by incorrectly adjusting the platform.

The simulated aircraft is outfitted with a baro-aided INS, a single GPS antenna, and an eight-channel GPS receiver. It is assumed there is no loss of GPS code-lock, no antenna shading, and no correlation between measurements noises. The model used for the INS is a generic, simplified model which is described in detail in Appendix A. The eight channels of the GPS receiver are looking at each of the four satellites twice. The differencing of redundant measurements taken from a satellite gives direct observability of measurement noise. This technique is described in Chapter 3.

It is important to note this research does not simulate many known, and often modelled, errors in the GPS measurement. These errors are outlined in Appendix B and are left for future research to include. The simulation is that of a benign flight path with no dynamic maneuvers. The INS is assumed to be functional at all times during the GPS integrity monitoring. This may at first seem an unrealistic assumption, but most military and commercial navigation systems require redundant INSs. This hardware redundancy prevents the declaration of a GPS failure when in actuality the INS failed. Instead, if an INS failure is detected, the GPS monitoring is no longer feasible. If an INS failure is isolated, there is still good data available to monitor the GPS. The application of errors involves some assumptions. These will be described in detail in Chapter 3 and 4.

1.5 Overview

This thesis is divided into five chapters and two appendices. Chapter 2 describes some theory and applications of Kalman filtering. It also contains a short section on chi-square hypothesis testing. Chapter 3 develops the differencing techniques used in the noise variance estimation algorithm. The chapter continues by describing the parallel filter structure used to detect failures (offsets) in individual GPS measurements. Then Monte Carlo simulation results of both the variance estimator and offset detection algorithm are analyzed in Chapter 4. The simulations are broken down into case studies that examine different scenarios of time-varying measurement noise variances and satellite failures. Chapter 5 summarizes the results and also offers recommendations for further research in this area. Appendix A contains the Kalman filter models used in the simulations shown in Chapter 4. Finally, Appendix B discusses GPS measurement errors and suggests models for use in future simulations.

Chapter 2 - Theory

2.1 Overview

Deterministic analysis is used to solve many problems in today's world, but for a large class of systems this type of analysis is inadequate. For these problems, the uncertainties in the models are significant factors in the behavior of the system. There are two areas where uncertainties typically enter into the model. First, dynamic models only approximate the true characteristics of a system. This approximation adds uncertainty to the adequacy of the dynamics equations. Furthermore, disturbances from the real world drive these dynamics equations, and these disturbances can be described only with some uncertainty as well. Second, sensors which measure system, or state, variables do not provide perfect information. Systems which use sensor output as feedback in a control loop are strongly affected by this type of measurement noise uncertainty. Methods have been developed to incorporate these ambiguities into the system model. The dynamics, measurements, and inputs can be modelled as stochastic processes [10, 14].

One widely used method is Kalman filtering. The following discussion outlines the specifics of Kalman filtering and the extensions used which allow the theory to be applied to a more general class of problems. A thorough development can be found in Maybeck [10]. The end result is that stochastic models better represent the real world and are absolutely essential in developing robust estimators and controllers.

A problem with embedding dynamics and measurement models in the filter arises when these models become incorrect, due to changes in the real world. There is no inherent adaptation or error detection executed in an ordinary KF. This adaptation and detection must be performed via some other route. It is common to construct a parallel algorithm which gives adaptive characteristics or error detection abilities [1, 11]. The algorithm presented in this work uses a parallel KF structure to attain the desired capabilities. The basic theory of this idea is described in Section 2.4. The last section in Chapter 2 contains a discussion of chi-square hypothesis testing. This theory will be used as a basis for error detection in this research.

2.2 The Kalman Filter

Kalman filtering can be categorized generally as an optimal linear estimation technique. These techniques have some method of fusing all available measurements to provide an 'optimal' estimate according to some measure of optimality. The advantage KFs have over most other techniques is the inclusion of *a priori* knowledge of the system which is producing the processes under observation. This advantage can quickly turn into a disadvantage if not properly implemented, i.e., if the assumed models are inadequate due to a change in the real world. A KF which assumes an incorrect dynamics model or measurement model can produce very *non*-optimal results. Implemented correctly, a KF outperforms many other algorithms by starting with a *propagated* estimate from an earlier time and using incoming measurements to refine, or *update*, that estimate. This propagate-update cycle is why the KF is referred to as a recursive algorithm. One important point is that the algorithm needs an initial condition from which to begin this cycle.

The optimal nature of the Kalman filter's estimates is based upon the assumptions on which the KF stands. The system modeled in the filter must be adequately described as linear and the noises which drive the system uncertainties are adequately portrayed as white, Gaussian processes. Whiteness implies the process has equal power density over all frequencies, and is a common simplification used by engineers to describe wide-band noise limited only by the system bandwidth. A linear shaping filter can be used to form a white noise process into one with desired time-correlated properties. This modeling technique is used to emulate various narrow band noises which may be present in a system. A Gauss-Markov process is one which is described completely by its mean, covariance, and covariance kernel. This is generally sufficient in most problems and can be shown to have physical justification as well. The assumption of linearity holds in numerous real-world problems, but cannot be applied to all cases. The extended Kalman filter (EKF) can handle a weakened assumption of linearity and is described later in this chapter.

2.2.1 Dynamics Model

A linear state model can be expressed as a stochastic differential equation in state space form [10]:

$$\dot{\mathbf{x}}(t) = \mathbf{F}(t)\mathbf{x}(t) + \mathbf{B}(t)\mathbf{u}(t) + \mathbf{G}(t)\mathbf{w}(t) \tag{1}$$

$\mathbf{x}$	state vector (nx1)
$\mathbf{F}$	homogeneous state dynamics matrix (nxn)
$\mathbf{B}$	control input matrix (nxr)
$\mathbf{u}$	deterministic control input matrix (rx1)
$\mathbf{G}$	driving noise input matrix (nxs)
$\mathbf{w}$	white, Gaussian driving noise vector (sx1)

Lower case bold letters represent vectors, upper case bold letters represent matrices, and normal or italicized lower case letters represent scalar variables. The $\mathbf{B}$ and $\mathbf{u}$ terms are zero in this study and will not be shown again. This research assumes $\mathbf{G}$ to be an identity matrix with $s = n$, and $\mathbf{w}$ is an n-dimensional vector with zero entries if necessary. The statistics of $\mathbf{w}$ are described below:

$$E\{\mathbf{w}(t)\} = \mathbf{0} \tag{2a}$$

$$E\{\mathbf{w}(t)\mathbf{w}^T(t+\tau)\} = \mathbf{Q}(t)\delta(\tau) \tag{2b}$$

$E\{\bullet\}$	expected value operator
$\mathbf{Q}$	dynamics noise strength matrix
$\delta(\bullet)$	Dirac delta function

Usually in navigation applications the states are not carried as whole values, but rather as errors [2]. This, along with the other assumptions described above, changes Equation (1) into Equation (3), letting $\boldsymbol{\delta}\mathbf{x}(t)$ represent error states:

$$\boldsymbol{\delta}\dot{\mathbf{x}}(t) = \mathbf{F}(t)\boldsymbol{\delta}\mathbf{x}(t) + \mathbf{w}(t) \tag{3}$$

2.2.2 Measurement Model

Real-world problems generally contain a continuous dynamics process with an output that is measured by sensors. Sensor measurements are sampled at discrete times when the measurement is fused with the current filter estimate. Kalman filters require a linear measurement model of the form shown in Equation (4). Note that $\mathbf{v}$ is a discrete-time, zero-mean, white, Gaussian process.

$$\mathbf{z}(t_i) = \mathbf{H}(t_i)\mathbf{x}(t_i) + \mathbf{v}(t_i) \tag{4}$$

The statistics of $\mathbf{v}$ are described below:

$$E\{\mathbf{v}(t_i)\} = \mathbf{0} \tag{5a}$$

$$E\{\mathbf{v}(t_i)\mathbf{v}^T(t_j)\} = \mathbf{R}(t_i) \quad \text{for } i = j \tag{5b}$$

$$E\{\mathbf{v}(t_i)\mathbf{v}^T(t_j)\} = \mathbf{0} \quad \text{for } i \neq j \tag{5c}$$

- $\mathbf{z}$ measurement vector (mx1)
- $\mathbf{H}$ observation matrix (mxn)
- $\mathbf{v}$ measurement noise vector (mx1)
- $\mathbf{R}$ measurement noise covariance matrix (mxm)

The measurement equation is converted to use error states and discrete time notation below:

$$\delta\mathbf{z}(t_i) = \mathbf{H}(t_i)\delta\mathbf{x}(t_i) + \mathbf{v}(t_i) \tag{6}$$

2.2.3 Propagation Equations

As discussed earlier, the KF propagates its state and covariance estimates forward in time until the next measurement is taken. Although digital computers cannot possibly integrate a time-varying, continuous function perfectly, the latest numerical integration techniques perform adequately well [13]. Propagation of the state vector is similar to the idea of knowing where you

should be before you get there, given the information at the current time, or state, and knowledge about the system's behavior. Propagation of the covariance matrix is akin to saying how accurately you should know where you will be before you get there. If that seems confusing, Equations (7) and (10) should make things clearer [10]. The error notation will be dropped so that the equations in this section remain true for the general case.

$$\frac{d}{dt}\hat{\mathbf{x}}(t/t_{i-1}) = \mathbf{F}(t)\hat{\mathbf{x}}(t/t_{i-1}) + \mathbf{B}(t)\mathbf{u}(t) \tag{7}$$

$$\frac{d}{dt}\hat{\mathbf{P}}(t/t_{i-1}) = \mathbf{F}(t)\hat{\mathbf{P}}(t/t_{i-1}) + \hat{\mathbf{P}}(t/t_{i-1})\mathbf{F}^T(t) + \mathbf{G}(t)\mathbf{Q}(t)\mathbf{G}^T(t) \tag{8}$$

The t/t_{i-1} term implies that the integration from t_{i-1} to t will be based on the measurement history through sample time t_{i-1}. Even though the filter-computed covariance, $\hat{\mathbf{P}}$, is estimated, designated by the hat, it is usually written as simply, $\mathbf{P}$.

2.2.4 Update Equations

It is conceivable that a sensor continuously outputs data to the filter, but due to the digital nature of most data processors, the normal implementation of a Kalman filter uses a discrete-time, or sampled-data, measurement model. That means the sensors are sampled at specified times when update calculations are performed. These updates are the optimal combination of incoming measurements, $\mathbf{z}(t_i)$, and present knowledge[2], $\hat{\mathbf{x}}(t_i^-)$ and $\mathbf{P}(t_i^-)$, to form new state and covariance estimates, $\hat{\mathbf{x}}(t_i^+)$ and $\mathbf{P}(t_i^+)$.

Kalman filtering incorporates sensor data into its estimates by determining a Kalman gain matrix, $\mathbf{K}$, based upon the current covariance matrix and the measurement model [10]:

$$\mathbf{K}(t_i) = \mathbf{P}(t_i^-)\mathbf{H}^T(t_i)[\mathbf{H}(t_i)\mathbf{P}(t_i^-)\mathbf{H}^T(t_i) + \mathbf{R}(t_i)]^{-1} \tag{9}$$

[2]The result of propagating Equations (7) and (8) to sample time t_i.

The updated state estimate is found by summing the current state with the contribution from the incoming measurements. The measurement contribution is the difference between the incoming measurement and the expected measurement premultiplied by the Kalman gain matrix:

$$\hat{\mathbf{x}}(t_i^+) = \hat{\mathbf{x}}(t_i^-) + \mathbf{K}(t_i)[\mathbf{z}_i - \mathbf{H}(t_i)\hat{\mathbf{x}}(t_i^-)] \tag{10}$$

The difference shown in the brackets of Equation (10) and explicitly in Equation (11) is often called the residual, denoted by $\mathbf{r}(t_i)$:

$$\mathbf{r}(t_i) = \mathbf{z}(t_i) - \mathbf{H}(t_i)\hat{\mathbf{x}}(t_i^-) \tag{11}$$

This value is extremely important in hypothesis testing and failure detection. The residual information is exploited in this work, as seen in Chapter 3 and described further in Section 2.5. The covariance update is shown below in Equation (12). Notice that the update can only reduce the covariance estimate; nevertheless, once $\mathbf{P}$ becomes positive definite, it will remain so [10].

$$\mathbf{P}(t_i^+) = \mathbf{P}(t_i^-) - \mathbf{K}(t_i)\mathbf{H}(t_i)\mathbf{P}(t_i^-) \tag{12}$$

2.3 Extended Kalman Filters

This section describes the methods used to apply the Kalman filter theory to a class of problems which exhibit nonlinear characteristics. It should be stated in advance that 'extended' Kalman filters (EKFs) cannot handle all nonlinear problems, only those whose perturbation from some (iteratively redeclared) 'nominal' can be well described by a first-order approximation over 'short' periods of time. This characteristic of EKFs should give the reader some insight into how they are implemented. Instead of a state vector multiplied by a dynamics matrix as in Equation (1), a nonlinear vector function, $\mathbf{f}$, is included in the differential stochastic equation [11]:

$$\dot{\mathbf{x}}(t) = \mathbf{f}[\mathbf{x}(t), \mathbf{u}(t), t] + \mathbf{G}(t)\mathbf{w}(t) \tag{13}$$

The $\mathbf{G}(t)$ term is assumed identity for all time and the white noise term has the same statistics as before, shown in Equation (2). In general, $\mathbf{x}$ is not a Gaussian process, although, for the work in this thesis, $\mathbf{f}$ is a linear function.

The discrete-time measurement model in extended Kalman filtering is generalized in a similar manner. The $\mathbf{Hx}$ term is replaced by $\mathbf{h}$, a nonlinear vector function [11]:

$$\mathbf{z}(t_i) = \mathbf{h}[\mathbf{x}(t_i), t_i] + \mathbf{v}(t_i) \tag{14}$$

Again, the statistics of $\mathbf{v}$ remain unchanged from Equation (17). These equations provide a representation of good nonlinear models, but this development has not yet yielded a tractable full-scale estimator solution.

The nonlinear equations can be linearized about a nominal trajectory so that the standard Kalman filter equations can be used. The extended Kalman filter *relinearizes* about a new state trajectory after every update cycle to propagate forward in time. The new trajectory used is based on the filter's best state estimate. The EKF propagation equations are shown here [11]:

$$\frac{d}{dt}\hat{\mathbf{x}}(t/t_{i-1}) = \mathbf{f}[\hat{\mathbf{x}}(t/t_{i-1}), \mathbf{u}(t), t] \tag{15}$$

$$\dot{\mathbf{P}}(t/t_i) = \mathbf{F}[t; \hat{\mathbf{x}}(t/t_i)]\mathbf{P}(t/t_i) + \mathbf{P}(t/t_i)\mathbf{F}^T[t; \hat{\mathbf{x}}(t/t_i)] + \mathbf{G}(t)\mathbf{Q}(t)\mathbf{G}(t)^T \tag{16}$$

The initial conditions of these differential equations after the filter's first propagation are given by:

$$\hat{\mathbf{x}}(t_i/t_i) = \hat{\mathbf{x}}(t_i^+) \tag{17a}$$

$$\mathbf{P}(t_i/t_i) = \mathbf{P}(t_i^+) \tag{17b}$$

The dynamics partial derivative matrix is defined by:

$$\mathbf{F}[t; \hat{\mathbf{x}}(t/t_i)] \triangleq \frac{\partial \mathbf{f}[\mathbf{x}, t]}{\partial \mathbf{x}}|_{\mathbf{x}=\hat{\mathbf{x}}(t/t_i)} \tag{18}$$

Similarly, after the propagation cycle, it *relinearizes* about the propagated state estimate to form a new observation matrix, $\mathbf{H}$, for use in the EKF update equations [11]:

$$\mathbf{H}[t_i; \hat{\mathbf{x}}(t_i^-)] \triangleq \frac{\partial \mathbf{h}[\mathbf{x}, t_i]}{\partial \mathbf{x}}|_{\mathbf{x}=\hat{\mathbf{x}}(t_i^-)} \tag{19}$$

$$\mathbf{K}(t_i) = \mathbf{P}(t_i^-)\mathbf{H}^T[t_i; \hat{\mathbf{x}}(t_i^-)]\{\mathbf{H}[t_i; \hat{\mathbf{x}}(t_i^-)]\mathbf{P}(t_i^-)\mathbf{H}^T[t_i; \hat{\mathbf{x}}(t_i^-)] + \mathbf{R}(t_i)\}^{-1} \tag{20}$$

$$\hat{\mathbf{x}}(t_i^+) = \hat{\mathbf{x}}(t_i^-) + \mathbf{K}(t_i)[\mathbf{z}_i - \mathbf{h}[\hat{\mathbf{x}}(t_i^-), t_i] \tag{21}$$

$$\mathbf{P}(t_i^+) = \mathbf{P}(t_i^-) - \mathbf{K}(t_i)\mathbf{H}[t_i; \hat{\mathbf{x}}(t_i^-)]\mathbf{P}(t_i^-) \tag{22}$$

Note the generalized residual vector premultiplied by $\mathbf{K}(t_i)$ in Equation (21).

2.4 Parallel Kalman Filter Architecture

The residuals of a KF can be monitored using various techniques to measure the 'correctness' of the model, or hypothesis, upon which the filter is based. When the hypothesis of a *single* filter is deemed 'bad' by some measure, it only knows that the models contained in the filter are incorrect to an unsatisfactory level. The filter doesn't inherently know how its model, or hypothesis, is bad, just that it is inadequate. This lack of knowledge could lead to a wandering or erratic search for the correct hypothesis. A parallel structure with *multiple* filters is a much more extensive and

proficient technique to detect changes of parameters or identify system failures. Each filter can have a different hypothesis, all of which are tested simultaneously. The output of the filter with the most correct model can be used for operational purposes as the other filters continue to be monitored to detect changes. This parallel type of structure can more quickly and accurately detect and identify changes in a system than a single filter.

One application of a parallel structure is the distributed Kalman filter (DKF). In general, a DKF is a set of filters using different information to estimate the same process. It is similar to asking several people standing on different street corners to describe a car accident that just happened. They have dissimilar views and will each describe the occurrence differently. After the accident, one of two things can happen. First, a police officer shows up on the scene to collect all the information from the individual witnesses and makes a decision, or estimate, based on all this information. Or second, the witnesses discuss the event among themselves, determine who had the best view of the scene, and agree to that version of the story. These two cases are analogous in the case of parallel Kalman filtering to the question of whether a master filter (the police officer) is used to fuse all the information of the elemental filters, or some type of logic is implemented to choose a single filter and reset the other filters to its estimates (discussion and consensus). The DKF generally uses a master filter to fuse information from elemental filters [3]. This research uses the second approach. The logic used to decide which filter has the most correct hypothesis is based upon chi-square testing, explained in Section 2.5. In this work, at certain times, measurements are fed into more than one filter. Trusted measurements are used in all the filters for all time. This creates a detection environment composed of a very tailored set of hypotheses. Also, the elemental filters output optimal estimates for each hypothesis and do not require a master filter to fuse the individual estimates.

2.5 Chi-Square Hypothesis Testing

Residuals are formed in the update process of a Kalman filter. This residual is a vector of length m, the number of measurements. The filter calculates the expected covariance of this residual

vector during the gain calculation. The filter-computed residual covariance, $\mathbf{A}$, is the bracketed term in Equation (9) and shown explicitly below:

$$\mathbf{A}(t_i) = \mathbf{H}(t_i)\mathbf{P}(t_i^-)\mathbf{H}^T(t_i) + \mathbf{R}(t_i) \tag{23}$$

If the model assumed in the filter adequately describes reality, the true residual covariance is equal to the computed covariance, $\mathbf{A}$. A summed chi-square random variable, $\chi(t_i)$, provides a test statistic that puts a quadratic penalty on variations of the residual vector over the last N samples:

$$\chi(t_i) = \sum_{j=i-N_\chi+1}^{i} \mathbf{r}^T(t_j)\mathbf{A}(t_j)^{-1}\mathbf{r}(t_j) \tag{24}$$

The $\mathbf{r}^T(t_j)\mathbf{A}(t_j)^{-1}\mathbf{r}(t_j)$ term in Equation (24) should be approximately m if the model (hypothesis) is good. If the model is bad, the residuals have larger mean-squared values than anticipated by the filter-computed $\mathbf{A}(t_j)$.

A threshold can be set empirically to detect a certain degree of failure or model inaccuracy. Said simply, if the summed chi-square variable is greater than a pre-specified threshold, a failure is declared. It is normally assumed that because chi-square testing only provides a single value every time sample, it does not have isolation ability. This is not altogether correct. It should be easily understood how multiple filters each performing a chi-square test would have the ability to isolate a failure. A chi-square test can be used to isolate failures in a multiple filter structure as well as in a *single* filter. A separate chi-square test can be executed on each scalar component of the measurement that is used to update the filter. In this fashion, test statistics are gathered on multiple scalar measurements and can be used to detect and isolate a failure of an individual sensor. Equations (25), (26), and (27) show the calculations used to form this scalar residual chi-square testing, letting $R_k(t_i)$ denote the k-th diagonal term of $\mathbf{R}(t_i)$, and $\mathbf{h}_k^T(t_i)$ denote the k-th *row* of the matrix $\mathbf{H}(t_i)$:

$$A_k(t_i) = \mathbf{h}_k^T(t_i)\mathbf{P}_k(t_i^-)\mathbf{h}_k(t_i) + R_k(t_i) \tag{25}$$

$$r_k = z_k(t_i) - \mathbf{h}_k^T(t_i)\hat{\mathbf{x}}(t_i^-) \tag{26}$$

$$\chi_k(t_i) = \sum_{j=i-N_\chi+1}^{i} \frac{r_k^2(t_j)}{A_k(tj)} \tag{27}$$

Equation (27) should yield approximately N_χ for a good hypothesis and much larger for a bad one. Using a larger N_χ reduces the probability of a false alarm being declared due to a single large measurement noise sample, but causes a lag in failure detection.

2.6 Summary

This chapter has described the general theory behind Kalman filtering, extended Kalman filtering, parallel filtering, and chi-square hypothesis testing. These are the tools exploited in Chapter 3 to provide protection from offsets in GPS measurements. This protection is strengthened by adding an adaptive measurement noise variance tuning algorithm described in the first sections of the next chapter.

Chapter 3 - Methodology

3.1 Overview

Chapter 3 focuses on the two main ideas of this thesis: noise variance estimation and signal offset error detection. It is stressed that these two algorithms are *completely independent* from each other. Either the noise estimation or the error detection algorithm can be performed alone. Moreover, due to the form of variance estimator used, it will not erroneously respond to a signal offset (such as a bias) in a real world scenario. Thus, the algorithm's performance is *decoupled.* The work shown here focuses on the error detection algorithm and aids this process by adaptively tuning the filters with information provided by the noise variance estimation. This creates a powerful environment in which to detect, isolate, and estimate changing parameters.

Section 3.2 describes the structure of the combined algorithm, as it is applied in this work. Section 3.3 discusses some general differencing techniques. Then Section 3.4 shows how differencing can be used to obtain direct observability of the measurement noise variance, without any coupling to satellite-dependent time-correlated errors. Next, Section 3.5 fully specifies the Kalman filter model used in this study. Finally, Section 3.6 presents many of the ideas behind the offset error detection algorithm.

3.2 Algorithm Structure

The structure proposed by this work to estimate errors and parameters in GPS measurements consists of three main components: a noise strength estimator, a parallel KF bank, and chi-square formulator. A block diagram of these three components, their associated logic, and system devices is shown in Figure 1.

First, the measurement noise variance is estimated by calculating statistics on a moving window of data. The KF is tuned real-time to adapt to real-world measurement noise variance changes. This process is thoroughly described in Section 3.4. The block diagram shows that the estimated

variance is not necessarily applied directly to the parallel set of filters. The algorithm allows for a modified value, R_f, to be placed in the KF models. This idea is discussed further in Section 3.6.5.

Second, the parallel KF bank, consisting of four filters, outputs navigation and detection data. The first filter of the bank is the navigation filter while the last three are detection filters. The navigation filter is solely responsible to output the best estimate of the errors in the INS. These estimates are applied as corrections to the output of the INS in real-time. The detection filters are never directly used to correct the INS, nor are they intended to provide precise state estimation. Their only purpose is to provide a solution with the benefit of one less GPS measurement. A comparison between the solution formed with the use of four satellites and one formed with only three provides a measure of the agreement of a single satellite with the rest of the system. If information from this one satellite varies significantly from the combined solution, the satellite is declared failed. The error detection concepts are covered in detail in Section 3.6.

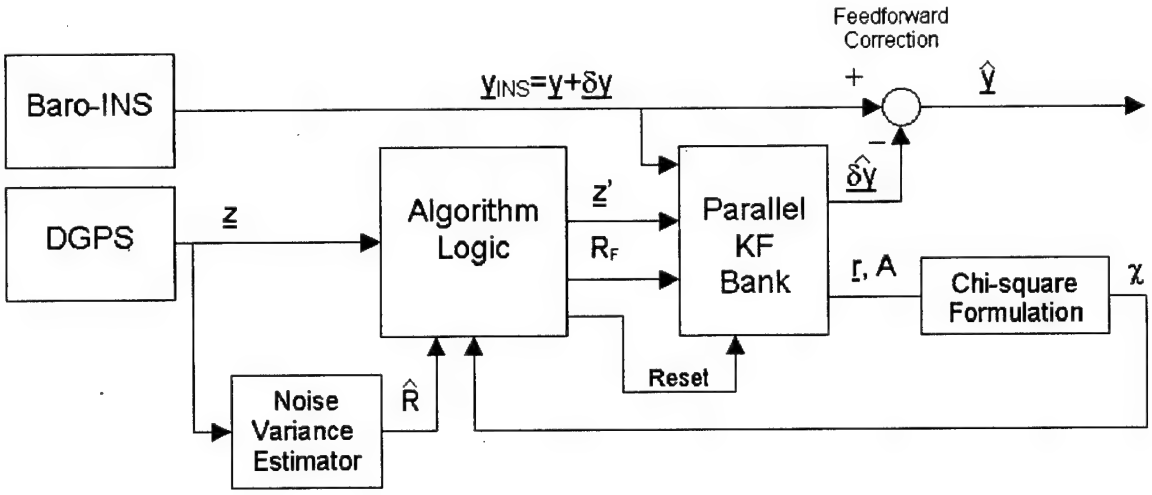

Figure 1. **Block Diagram of Combined Algorithm**

The third component in the algorithm uses the filter residuals to form chi-square variables after every measurement update. This component is just a nonlinear vector function which returns a vector of chi-square variables from the detection filters. This vector is used in the algorithm logic to decide whether a failure has occurred.

The algorithm logic block in Figure 1 is used to designate the decision making components of the algorithm. The operation of this block, including the critical filter resets, is fully described in Section 3.6. Notice the measurements applied, $\mathbf{z}'$, are not necessarily identical to the available measurements, $\mathbf{z}$. This aspect of the algorithm emphasizes that each filter can receive a different measurement set and these sets are allowed to change if needed.

3.3 Standard Differencing Techniques

Equation (28), explained in detail in Appendix A, shows the model for a pseudorange (PR) measurement. [3]

$$\rho_k^i = R_k^i + \delta t_k + \delta t^i + I_k^i + T_k^i + M_k^i + C_k^i + v_k^i \tag{28}$$

ρ_k^i	Pseudorange of satellite i calculated by receiver k
R_k^i	True range between satellite i and antenna k
δt_k	Receiver k clock range error
δt^i	Satellite i clock range error
I_k^i	Ionospheric delay between satellite i and antenna k
T_k^i	Tropospheric delay between satellite i and antenna k
M_k^i	Multipath delay between satellite i and antenna k
C_k^i	Dynamics-induced code tracking loop error of receiver k on the i^{th}satellite signal
v_k^i	White noise error of receiver k on the i^{th} satellite signal

This equation shows all the errors associated with the GPS range measurement: seven modelled errors and one lumped, unmodeled error. The error designated by v is the lumped error, and it is assumed to be a zero-mean, white, Gaussian process. *All the errors have units of distance* even though some originate from timing errors. Hence, they are referred to as *range* errors.

[3]Note: the δ symbol is used to denote an error or error state.

3.3.1 Single Differencing (Two-Receiver)

Normal differencing, i.e. differential GPS (DGPS), uses differences of pseudorange measurements from two receivers to one satellite. This eliminates all satellite dependent errors, δt^i, and significantly reduces highly-correlated atmospheric errors, I_k^i and T_k^i [16]. Although this increases uncorrelated errors, M_k^i, C_k^i, and v_k^i, these are generally much smaller than the clock and atmospheric errors. All these factors make DGPS a highly effective and simple method to improve positioning accuracy. Two-receiver differencing (or DGPS) is outlined in this section.

3.3.1.1 Ephemeris Error. The differential correction is the difference of the measured range, PR, and the expected range to a reference station. The pseudoranges are output from the GPS receiver and the expected range is calculated below:

$$\hat{R}_{ref}^i = \left|\hat{\mathbf{r}}_{sv}^i - \mathbf{r}_{ref}\right| \tag{29}$$

$\hat{R}_{ref}^i$ Expected range from reference station to SV i
$\hat{\mathbf{r}}_{sv}^i$ Position of SV i according to the ephemeris
$\mathbf{r}_{ref}$ Surveyed position of reference station

This equation can be used only when the user's position is known, as it is in the case of a fixed reference station. The only problem is that the SV's position is not known to the desired precision. This error, or uncertainty, is referred to as ephemeris error. The three-dimensional ephemeris error can be represented as a vector in Cartesian coordinates, $\boldsymbol{\delta}\mathbf{r}_{sv}^i$. This error is projected onto the unit vector in the direction of the signal path from the reference to satellite i, $\bar{\mathbf{a}}_r^i$, to form the one-dimensional effect, δE_r^i, on the pseudorange. Notice the ref subscript is shortened to just r. This projection is performed through an inner product operation, denoted by the dot operator $(\cdot)$:

$$\delta E_r^i \triangleq \bar{\mathbf{a}}_r^i \cdot \boldsymbol{\delta}\mathbf{r}_{sv}^i \tag{30}$$

3.3.1.2 Differential Correction. To form the differential correction, the expected range, $\hat{R}_r^i$, is calculated as the true range, R_r^i, minus the effect of the ephemeris error:

$$\hat{R}_r^i = R_r^i - \delta E_r^i \tag{31}$$

Now, the difference between the measured PR and this expected range forms the differential correction:

$$\delta\rho^i = \rho_r^i - \hat{R}_r^i \tag{32}$$

This reduces to the sum of all errors at the reference station:

$$\delta\rho^i = (R_r^i + \delta t_r + \delta t^i + I_r^i + T_r^i + M_r^i + C_r^i + v_r^i) - (R_r^i - \delta E_r^i) \tag{33a}$$

$$= \delta E_r^i + \delta t_r + \delta t^i + I_r^i + T_r^i + M_r^i + C_r^i + v_r^i \tag{33b}$$

The differential corrections for all the satellites in view are broadcast to all the users in the area. The differential corrections are received and applied to the appropriate PR. This effectively achieves the single difference between the user's receiver and the fixed receiver at the known reference point. This idea is shown in mathematical form below:

$$\begin{aligned}\Delta\rho_{ur}^i &= \rho_u^i - \delta\rho^i \\ &= (R_u^i - \delta E_r^i) + (\delta t_u - \delta t_r) + (I_u^i - I_r^i) \\ &+ (T_u^i - T_r^i) + (M_u^i - M_r^i) + (C_u^i - C_r^i) + (v_u^i - v_r^i)\end{aligned} \tag{34}$$

It should be noted that the single differenced measurement is applied with the same observation matrix, H, as the raw pseudorange.

3.3.1.3 Effects of Differencing. When the differential correction is subtracted from the PR calculated at the user receiver, several changes occur:

(1) Satellite clock error, δt^i, is removed

(2) Ephemeris error is subtracted[4]

(3) Ionospheric error is reduced

(4) Tropospheric error is reduced

(5) Multipath error is increased[5]

(6) Code loop error is increased

(7) Measurement noise is increased

These changes nearly always significantly improve the position solution. The amount of reduction of the ionospheric and tropospheric errors depend on the correlation between the user and reference errors. Both errors are temporally and spatially correlated. When the user is far from the reference, the improvement is less. When the correction is applied to a time for which it is not valid, the improvement is less. Both atmospheric error correlations degrade gracefully with time and distance [15].

Differential corrections can be applied by three methods. First, the corrections can be stored in a file and applied to the user's stored navigation data. This provides a post-processing approach very useful at test ranges. Second, the user's navigation data can be transmitted back to a reference station which applies the corrections. This method is used when there are many standard GPS users who have no local need for positioning information, but are tracked precisely by a central station. Such users include city emergency vehicles and commercial trucking agencies. Lastly, the corrections can be broadcast to the user and applied real-time. This approach is used by the

[4]If the differenced measurement is applied as an error state measurement, the estimated range to the user is subtracted from the measurement. This will cause the true range and the ephemeris error to be removed. This assumes the line of sight vectors from the user to the SV and from the reference to the SV are in the same direction. In reality, there is a very small angular difference, a maximum of 1.4° per 500 km. The effect is in the cm range and is usually neglected.

[5]It can be assumed that the reference point is located such that no large multipath risks are nearby. Also, narrow band correlators and choke ring antennas should be used to reduce multipath measurement errors as much as possible. The multipath addition from the reference station should be very small.

military in operational situations that require real-time high precision navigation solutions. It is also the approach implemented in this research.

3.3.2 Double Differencing (Two-Receiver, Two-Satellite)

Double differencing is usually implemented as the difference of two DGPS measurements to two satellites. This eliminates user clock error altogether, but also reduces the number of measurements to $n-1$, n being the number of GPS satellites available. By differencing two single differenced measurements, as in Equation (34), a double differenced measurement is formed:

$$\begin{aligned}\nabla\Delta\rho_{kl}^{ij} &= \Delta\rho_{kl}^{i} - \rho_{kl}^{j} \\ &= (R_k^i - R_l^i) - (R_k^j - R_l^j) + (I_k^i - I_l^i) - (I_k^j - I_l^j) \\ &+ (T_k^i - T_l^i) - (T_k^j - T_l^j) + (M_k^i - M_l^i) - (M_k^j - M_l^j) \\ &+ (C_k^i - C_l^i) - (C_k^j - C_l^j) + (v_k^i - v_l^i) - (v_k^j - v_l^j) \end{aligned} \tag{35}$$

Equation (35) shows the double difference between receivers k and l, and satellites i and j. This method is usually only implemented for carrier phase ambiguity resolution and not used for code measurements, since it does not significantly improve the position solution. The specifics of double differencing measurement applications are not shown in this work.

3.3.3 Triple Differencing (Two-Receiver, Two-Satellite, Two-Epoch)

Taking one more step, a difference of two double differenced measurements at two times, or epochs, forms what is commonly referred to as a triple difference. This method is normally used only for static implementations, such as surveying. Since its performance degrades quickly in a dynamic environment, it is not used in this study.

3.4 Measurement Noise Variance Estimation

The estimation process is based on a statistical evaluation of the differences of redundant measurements over a specified period of time. It is important to note that this variance is associated

with the measurement process and has units of distance squared. The measurement noise is assumed to be a normally distributed, white, zero-mean process. Over the period of evaluation, it is assumed to be strictly stationary and ergodic. It may seem ironic that the process is assumed to have time-invariant statistical characteristics when the algorithm's purpose is to track changes in the statistical properties of a process. This issue is resolved by only evaluating the process sample variance over a prespecified number of samples.

3.4.1 Redundant Measurement Differencing (Two-Channel)

The redundant measurement differencing (RMD) method presented in this work forms the difference between redundant measurements made by one receiver to one satellite. These measurements are taken in *two channels* within a single receiver, k:

$$\rho_k^i = R_k^i + \delta t_k + \delta t^i + I_k^i + T_k^i + M_k^i + C_k^i + v_k^i \tag{36a}$$

$$\rho_{k'}^i = R_{k'}^i + \delta t_{k'} + \delta t^i + I_{k'}^i + T_{k'}^i + M_{k'}^i + C_{k'}^i + v_{k'}^i \tag{36b}$$

The prime designates the second, or redundant, measurement taken by the same receiver k. This implies that the ranges and all errors, except one, are equal to those in the second measurement. In other words, the atmospheric errors in the measurements of two channels are equal because there is only one signal travelling from one satellite to one antenna, i.e. the signal path is identical. The two clock errors are equal by definition since the same clock is used to generate both measurements at the same time. The dynamics of both channels are equal and, therefore, so is the dynamics-correlated code loop error. The only non-identical error is the white noise term. Using no approximations, the redundant measurement difference reduces to:

$$\Delta\rho_{kk'}^i = \rho_k^i - \rho_{k'}^i = (v_k^i - v_{k'}^i) \tag{37}$$

3.4.2 Moment Calculations

It is desired to calculate the statistical moments of the white measurement noise for their use in the KF models. An empirical mean of v need not be found, since it is easily shown that the true mean of $\Delta\rho^i_{kk'}$ is zero:

$$E(\Delta\rho^i_{kk'}) = E(v^i_k - v^i_{k'}) = E(v^i) - E(v^i) = 0 \tag{38}$$

Since k and k' denote two realizations of identical processes, the subscript is dropped for simplicity.

The second central moment, variance, of the difference measurement is approximated at the present time, t_p, by Equation (39). This calculation assumes the process is ergodic over the past N_R samples.

$$E\{(\Delta\rho^i_{kk'})^2\} = \frac{1}{(N_R - 1)} \left[\sum_{j=p-N_R+1}^{p} \left\{\Delta\rho^i_{kk'}(t_j)\right\}^2 \right] \tag{39}$$

If the expected value of v^i is assumed to be zero and if v^i_k and $v^i_{k'}$ are assumed to be independent, it can be shown that the variance of the difference measurement is twice the variance of an individual noise process[6]:

$$\begin{aligned} E\{(\Delta\rho^i_{kk'})^2\} &= E\{(v^i_k - v^i_{k'})^2\} \\ &= E\{(v^i)^2\} + E\{(v^i)^2\} - 2E\{v^i_k v^i_{k'}\} \\ &= 2E\{(v^i)^2\} \end{aligned} \tag{40}$$

This relationship can then be used to form the estimate of R^i at the present time t_p:

[6]The subscript j is a dummy variable of summation, the superscript i represents the satellite number, and the subscript p denotes the present epoch.

$$\widehat{R}^i(t_p) \quad = \quad E\{(v^i)^2\} = \frac{E\{(\Delta\rho^i_{kk'})^2\}}{2} \tag{41}$$

The value for N_R reflects how constant the noise statistics are thought to be. N_R could be adaptively changed to smaller values during high-dynamic maneuvers or when approaching areas where GPS signals are likely to be jammed. Also notice the division by two of the noise sample variance. This corrects for the measurement being a difference of two independent processes with equal variance. It is extremely important to understand that either the noise variance estimation algorithm or the error detection algorithm can be run together or separately. This estimate is *completely independent* of any offset or error the whole-valued measurement is experiencing; such an effect would be cancelled out within the differencing performed in Equation (37). This *decouples* the noise variance estimation from the signal offset detection problem.

3.5 Filter Model

3.5.1 Dynamics

This study uses a 13-state extended Kalman filter with a linear dynamics model and a nonlinear measurement model. This 13-state model is used for the navigation filter, the detection filters, and the truth model. The first nine states directly relate to the standard INS Pinson error states, i.e. three-dimensional position, tilt, and velocity. The next two states are dedicated to the barometric altimeter. These states are the altitude error above the reference ellipsoid and the total barometric altimeter time-correlated error. The last two states model the GPS measurement. User clock bias and drift are the only two errors estimated by the filter. The 13-dimensional state vector is shown below:

$$\delta\mathbf{x} \quad = \quad [\delta\theta^T \;\; \phi^T \;\; \delta\mathbf{V}^T \;\; \delta h \;\; \delta h_B \;\; \delta t \;\; \delta\dot{t}]^T \tag{42}$$

$\delta\theta$	Computer frame error angles
ϕ	Platform tilt angles
$\delta\mathbf{V}$	INS velocity errors
δh	Altitude error above the reference ellipsoid
δh_B	Total barometric altimeter time-correlated error
δt	GPS user clock phase error
$\delta\dot{t}$	GPS user clock frequency error

The dynamics matrix can be described as submatrices directly associated with the state vector partitions shown above. The overall matrix is given in Equation (43) and the individual submatrices are explained in detail in Appendix A.

$$\mathbf{F}_{(13\text{x}13)} = \begin{bmatrix} \mathbf{F11}_{(3\text{x}3)} & \mathbf{0} & \mathbf{F13}_{(3\text{x}3)} & \mathbf{0} & \mathbf{0} \\ \mathbf{F21}_{(3\text{x}3)} & \mathbf{F22}_{(3\text{x}3)} & \mathbf{F23}_{(3\text{x}3)} & \mathbf{0} & \mathbf{0} \\ \mathbf{F31}_{(3\text{x}3)} & \mathbf{F32}_{(3\text{x}3)} & \mathbf{F33}_{(3\text{x}3)} & \mathbf{F34}_{(3\text{x}2)} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{F43}_{(2\text{x}3)} & \mathbf{F44}_{(2\text{x}2)} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{F55}_{(2\text{x}2)} \end{bmatrix} \tag{43}$$

White noise terms are added to each state to form the full stochastic dynamics equation:

$$\delta\dot{\mathbf{x}}_{(13\text{x}1)} = \mathbf{F}_{(13\text{x}13)}\delta\mathbf{x}_{(13\text{x}1)} + \mathbf{w}_{(13\text{x}1)} \tag{44}$$

All the white noise processes are assumed to be zero-mean and independent; therefore, they are completely described by the thirteen diagonal strength terms. Specific values for the dynamics matrix and dynamics driving noise strength matrix are contained in Appendix A. Initial conditions do not need to be specificed since this simulation was performed at steady state.

3.5.2 Measurements

The measurement model incorporates one barometric altimeter altitude and four GPS pseudo-ranges into the filter estimates. The 5-dimensional measurement vector is shown here:

$$\delta\mathbf{z} = [\delta z_{baro} \;\; \delta z^1_{sv} \;\; \delta z^2_{sv} \;\; \delta z^3_{sv} \;\; \delta z^4_{sv}]^T \tag{45}$$

The barometric altimeter is used to generate a difference measurement between altimeter and INS indications of altitude to eliminate the whole state form of the variable. The INS altitude, h_{INS}, is modeled as the true altitude, h_{true}, plus the INS error in altitude above the reference ellipsoid, δh (state ten of the INS model). The altimeter output, h_{baro}, is modeled as the sum of the true altitude, the total barometric altimeter correlated error, δh_B (state eleven of the INS model), and a white noise term. This measurement is linear and the equations are shown explicitly below:

$$\delta z_{baro} = h_{INS} - h_{baro} \tag{46a}$$

$$= [h_{true} + \delta h] - [h_{true} + \delta h_B - v_{baro}] \tag{46b}$$

$$= \delta h - \delta h_B + v_{baro} \tag{46c}$$

$$= \delta x_{10} - \delta x_{11} + v_1 \tag{46d}$$

The DGPS measurement model follows directly from Section 3.3.1. Due to model reduction, the DGPS measurement, for this implementation, has significantly fewer terms:

$$\Delta\rho^i = R_u^i + (\delta t_u - \delta t_r) + (v_u^i - v_r^i) \tag{47a}$$

$$= R_u^i + \delta t + v^i \tag{47b}$$

Notice the simplification in the notation, i.e. the subscript ur on the difference measurement is dropped. The two clock errors are combined to form one state, δt (state twelve of the INS model). Similarly, the white noise terms are combined to form a single noise.

Since R_u^i, in Equation (47), is not known to the receiver, an expected range from the user position to each satellite is calculated. The magnitude of the difference between the satellite position, $\mathbf{r}_{sv}^i$, and the user position, $\mathbf{r}_u$, represents the expected range measurement:

$$\Delta\hat{\rho}^i = \left|\mathbf{r}_u - \mathbf{r}^i_{sv}\right| \tag{48a}$$

$$= \sqrt{(x_u - x^i_{sv})^2 + (y_u - y^i_{sv})^2 + (z_u - z^i_{sv})^2} \tag{48b}$$

The nonlinear form of Equation (48b) is unusable by a linear Kalman filter. This forces us to express it as a Taylor series expanded about the current INS position output and the SV position according to the ephemeris information, $\mathbf{r}_{INS}$ and $\hat{\mathbf{r}}^i_{sv}$ respectively. By truncating the series at first-order, a linear approximation is formed:

$$\begin{aligned}\Delta\hat{\rho}^i \;\cong\; & R^i_u + \left\{ \frac{\partial \Delta\hat{\rho}^i(\mathbf{r}_u, \mathbf{r}^i_{sv})}{\partial \mathbf{r}_u}\Big|_{\mathbf{r}_u = \mathbf{r}_{INS}; \mathbf{r}^i_{sv} = \hat{\mathbf{r}}^i_{sv}} \right\} \partial \mathbf{r}_u \\ & + \left\{ \frac{\partial \Delta\hat{\rho}^i(\mathbf{r}_u, \mathbf{r}^i_{sv})}{\partial \mathbf{r}^i_{sv}}\Big|_{\mathbf{r}_u = \mathbf{r}_{INS}; \mathbf{r}^i_{sv} = \hat{\mathbf{r}}^i_{sv}} \right\} \partial \mathbf{r}^i_{sv}\end{aligned} \tag{49}$$

Evaluating the partial derivatives in the above equation (finding that for *this model* $\partial \mathbf{r}^i_{sv} = 0$), yields the following:

$$\Delta\hat{\rho}^i \;=\; R^i_u - \left[\frac{\hat{x}^i_{sv} - x_{INS}}{|\mathbf{r}_{INS} - \hat{\mathbf{r}}^i_{sv}|}\right] \partial x_u - \left[\frac{\hat{y}^i_{sv} - y_{INS}}{|\mathbf{r}_{INS} - \hat{\mathbf{r}}^i_{sv}|}\right] \partial y_u - \left[\frac{\hat{z}^i_{sv} - z_{INS}}{|\mathbf{r}_u - \hat{\mathbf{r}}^i_{sv}|}\right] \partial z_u \tag{50}$$

Finally, the DGPS pseudorange difference measurement is formed:

$$\begin{aligned}\delta z^i_{sv} \;&=\; \Delta\rho^i - \Delta\hat{\rho}^i \\ &=\; \left[\frac{\hat{x}^i_{sv} - x_{INS}}{|\mathbf{r}_{INS} - \hat{\mathbf{r}}^i_{sv}|}\right] \partial x_u + \left[\frac{\hat{y}^i_{sv} - y_{INS}}{|\mathbf{r}_{INS} - \hat{\mathbf{r}}^i_{sv}|}\right] \partial y_u + \left[\frac{\hat{z}^i_{sv} - z_{INS}}{|\mathbf{r}_{INS} - \hat{\mathbf{r}}^i_{sv}|}\right] \partial z_u + \delta t + v^i\end{aligned} \tag{51}$$

The INS position and SV ephemeris position are known values which can be directly substituted into the measurement equation. The δt term is the twelfth state variable of the system. The three

user position errors can be calculated through a linear transformation of the first three states, since they are actually tilt errors. The mean and variance of v^i are set at 0 ft and 9 ft^2, respectively.

This simulation uses identical truth and filter model designs; therefore, only user clock phase and frequency errors are generated in the simulation. This reduced-order truth model is sufficient to demonstrate a proof-of-concept, as opposed to providing a full appraisal of performance to be anticipated in real-world operation. A full-scale GPS error model is described in Appendix B.

3.6 Offset Detection

3.6.1 Parallel Structure

As stressed throughout this thesis, the parallel structure is the key to the detection and isolation characteristics of this algorithm. Figure 2 is used to describe the form of the offset detection algorithm when testing a single satellite. This figure only shows three detection filters looking for offsets in one satellite. Simultaneous monitoring of four satellites would require twelve detection filters and one for navigation. Figure 2 breaks down the algorithm logic block referred to in a general sense in Figure 1. First, the GPS measurement vector is received and directed to the various filters. If a satellite has been declared failed, it is withheld from the navigation filter. Next, the navigation filter outputs its estimate to correct the indicated INS position. Meanwhile, the detection filters output their residuals, **r**, and filter-computed residual covariance, **A**, to the chi-square summation block. The chi-square variables are tested against an upper threshold at which a failure, or unacceptable offset, is declared. This declaration, represented by the dashed line, causes a single binary bit to be sent to the algorithm logic block which accomplishes two tasks: it informs the navigation filter to be reset to the values contained in the detection filter which has not been reset for the longest time (refer to Section 3.6.2), and it sets the figure of merit (FOM) of the failed satellite to zero. The FOM of a satellite is either one, fully functional, or zero, fully failed. If the FOM of a satellite is zero, that satellite is not used in the navigation KF.

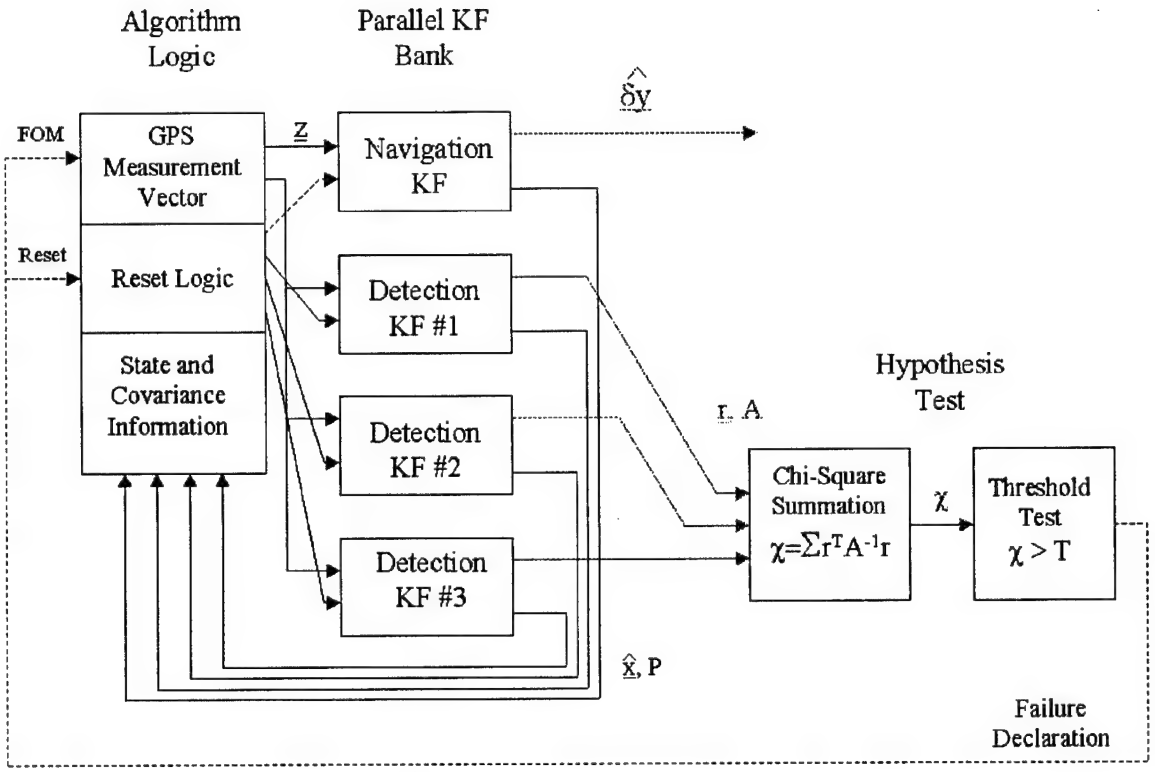

Figure 2. **Offset Detection Algorithm Structure**

3.6.2 Detection Filter Resets

The detection capability of this algorithm ultimately lies in the reset logic of the elemental Kalman filters. The navigation filter is updated with all GPS satellite measurements (four per update time) to estimate the errors in the baro-INS of Figure 1. The detection filters are propagated through time and updated with one fewer GPS satellite measurement, three, than the navigation filter. This is done intentionally to increase any discrepancy between the satellite which is not included and all the others. The measurements from the satellite under test are still brought into the filter and residuals are formed, but *no update is performed.* The residual information is evaluated for a prespecified amount of time and, if no failure is declared, the detection filter is reset to the state estimate and covariance matrix contained in the navigation filter, thereby incorporating the benefit of *all four* GPS measurement time histories up to that time. This increases the accuracy of the estimates in the detection filter, and thus increases the ability to detect an error in a specific satellite. A set of detection filters is used for each hypothesis, i.e., correct operation of a single satellite. This allows a staggered resetting of the detection filters to provide continuous testing of that hypothesis. A graphical depiction of the reset timing of three detection filters and their associated $0\pm 1\sigma$ (standard deviation) bounds under a no-fail condition is shown in Figure 3.

The top plot represents the navigation filter being updated every second by all four GPS measurements. It shows the filter-computed $\pm 1\sigma$ of a scalar residual. In between each measurement epoch, the position estimate will drift slightly due to the INS characteristics. The next plot shows the residuals for detection filter #1. It runs for 18 seconds without updates from the satellite under test. If no failure is detected, it is reset to the navigation filter's estimate to gain the benefit of higher precision for later detection. The drop in the standard deviation is due to the *reset*, not a measurement update. The drift shown in the detection filter's plot represents the increase in the covariance values when the measurement set is reduced from four satellites to three [8]. The last two plots show the other two detection filters. These filters are acting similarly to the first detection filter, but each lead the previous filter by six seconds.

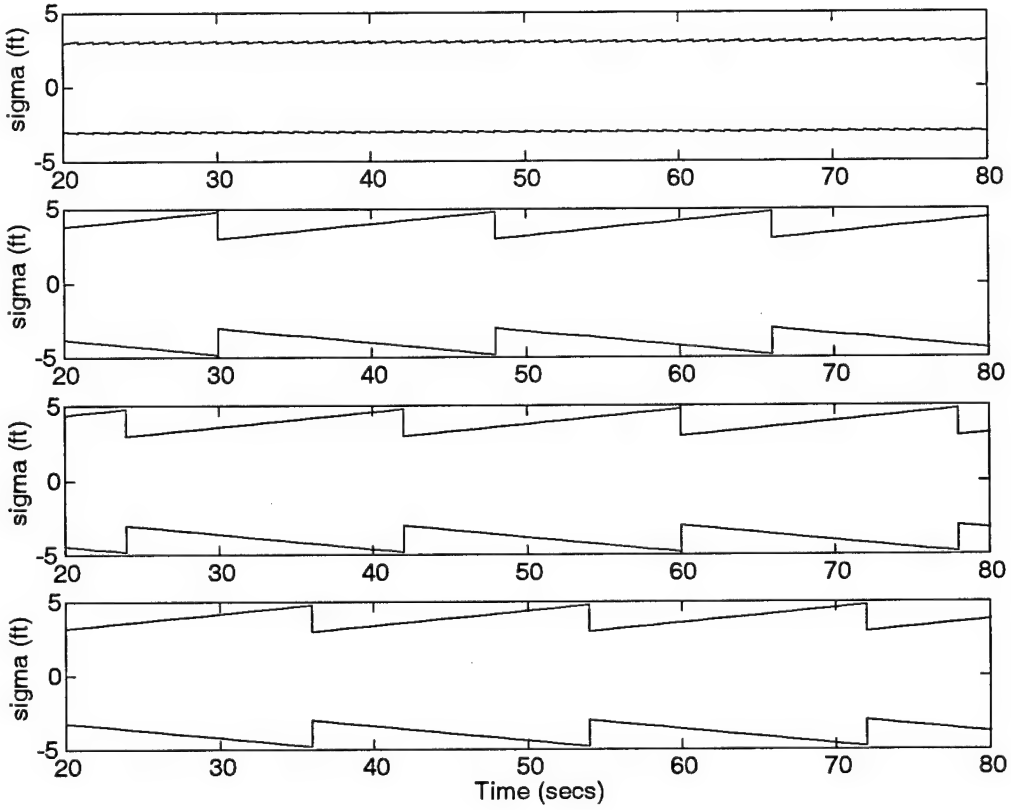

Figure 3. **Reset Timing Diagram: Nominal Operation**

3.6.3 Detection Limits

This section compares the detection sensitivity of two filters with different update rates through the use of a *simplified* example. Figure 3 highlights the reason why a filter which is not updated for a period of time is more sensitive to failures. Consider a system in which a failure is declared when the residuals continually break a 1σ boundary as shown in the uppermost plot of Figure 3. Let that one-sigma value be σ_o after an update. The system error increases between updates and, immediately before the update, the one-sigma value (denoted as σ) is equal to σ_o plus the system drift rate, λ_s, times the update period, T, shown below [7]:

$$\sigma = \sigma_o + \lambda_s T \tag{52}$$

Now, for a filter updated at a slower rate, say eighteen times more slowly, as for any of the three detection filter plots in Figure 3, the following value is attained:

$$\sigma = \sigma_o + \lambda_s(18T) \tag{53}$$

The detection of ramp offsets in the measurement signal is of primary concern in this research. The error drift rate (or offset ramp rate), λ_e, which will break the boundary for the fast filter is given by:

$$\lambda_{efast} = \frac{\sigma_o + \lambda_s T}{T} \tag{54}$$

Similarly for the slow filter:

$$\lambda_{eslow} = \frac{\sigma_o + \lambda_s(18T)}{18T} \tag{55}$$

[7] This assumes the error starts at zero and the offset error is a constant ramp.

By comparing Equations (55) and (56), it is obvious the slower filter is more sensitive to drifting or ramping failures. Consider σ_o=3 ft, λ_s=0.5 ft/sec, and T=1 sec. This yields detectable drift rates of 3.5 ft/sec and 0.667 ft/sec, respectively. These numbers demonstrate significant improvement and are representative of the GPS/INS problem, but this is not a thorough analysis. It serves only as a simple example to demonstrate the advantages of slower update rates. The theoretical limit to reliable error detection is fundamentally set at the system drift rate. It is foolish to think errors can be detected beneath the statistical variation of the system in question. In this case the system is a Baro-INS updated by *three* GPS measurements. The system drift remains nearly zero if a constant number of satellites in a static geometry are continually input to the filter. A temporary drift is generated when the number or geometry of satellites is changed.

3.6.4 Navigation Filter Resets

If an error *is declared*, the navigation filter is reset to the information in the detection filter whose residuals led to the failure declaration. At first, this resetting of the navigation filter to estimates achieved with one fewer satellite may seem to be a senseless degradation of the current navigation solution. Although this action increases the filter-computed covariance, it is intended to remove the error induced by a satellite which has most likely failed at some time before the actual declaration. This technique is especially effective when the signal error takes the form of a slowly increasing offset. Without such a reset method, even if a satellite is excluded from use in the filter after the time of declaration, the filter estimates have *already been corrupted* by that satellite.

The research in this thesis was first carried out by propagating the detection filters without the benefit of any GPS measurements. Then, all four satellite measurements were compared to the estimates of the detection filter, basically simulating a drifting INS. This method was found to be capable of detecting errors, but only of high magnitudes. Adding selected satellite measurements to aid the detection filters enhances the detection ability, though at the cost of narrowing the search space to one satellite per detection filter bank. Overall, this method effectively lowers the system drift rate mentioned before by including three satellites as part of the system.

3.6.5 Hypothesis Testing

Each set of three filters has a single hypothesis (that a single satellite under test has not failed). These filters are *never* updated by the measurements from that satellite; however, the filters do form residuals with this satellite. This is performed to create residuals conditioned on measurements only from satellites *not* under test. Each filter forms a chi-square variable, described in Chapter 2, through the use of its residuals. If this chi-square variable breaks a threshold, the hypothesis that the residual information is indicative of an *offset-free* measurement is declared false. After a failure is declared, the navigation filter is reset as described in Section 3.6.4 and the failed satellite is not used to update filters for the remaining duration of the simulation. Failed satellites are not brought back on-line in this simulation due to the increase in the variance of residuals after the loss of a satellite. If a chi-square lower threshold is set so satellites can be brought back on-line, a constant signal offset which had caused a failure declaration could become declared as 'good' as the residual variance bounds increase. This effect has been demonstrated by the research performed in this thesis, and it is strongly recommended for any further work to consider carefully the manner in which an algorithm re-incorporates failed satellites. One reasonable technique would be to select an unfailed satellite immediately from a set that is being monitored for incorporation into the solution as the fourth satellite after a failure. By quickly removing the failed satellite and adding a different one, the transient effect seen in residual covariance should last for only a short time. In this manner, the occurrence of a failure would cause no decrease in navigation performance or detection capability. Ideas of this nature are discussed more in depth in Chapter 5.

3.6.6 Adaptive Tuning

The noise estimation algorithm yields an $\hat{R}$ for each satellite. In an environment of quickly changing signal noise strength, the ergodic assumption used in Equation (39) fails to describe reality accurately. Thus, this assumption leads to a tracking lag in estimation of the current $\mathbf{R}$. If the estimate lags too greatly, an incorrect value, $\mathbf{R}_f$, appears in the filter and can cause the residuals to become statistically much larger than the filter-computed residual covariance when the *true* $\mathbf{R}$

is increasing. This may cause a false alarm failure declaration of a satellite by the detection logic looking for signal offsets, due to changing interference levels. Such false alarms should be avoided whenever possible. When the estimate lags when R is decreasing, the residuals are much smaller than expected and the algorithm is suspectable to missed alarms. The value of N_R can be reduced to attempt to decrease the lag in the estimate of the noise variance, but this provides a variance estimate of higher error variance. In addition to lowering N_R, the filter measurement covariance, $\mathbf{R}_f$, can be set to a value higher than the estimated covariance by the addition of a constant matrix, $\mathbf{R}_o$:

$$\mathbf{R}_f = \mathbf{R}_o + \hat{\mathbf{R}} \tag{56}$$

This second technique is a conservative tuning method which slightly decreases performance during nominal operation, but keeps the false alarm rate down. One possible implementation of this idea might be to add $\mathbf{R}_o$ only when a significant change in $\hat{\mathbf{R}}$ occurs. These two techniques can be employed in conjunction with each other or separately.

3.7 Summary

This chapter focuses on two extremely different, yet quite complementary, techniques to increase the confidence in the navigation solution of an integrated INS/GPS system. The first part of the chapter describes the redundant measurement differencing (RMD) method for estimating measurement noise variance and highlights the uniqueness of this method when compared to other standard differencing techniques. The second half is concerned with a failure detection algorithm for measurement signal offsets which is based directly on the parallel KF structure and reset logic. The performance of these two algorithms in simulation is summarized in Chapter 4.

Chapter 4 - Results and Analysis

4.1 Overview

This chapter is divided into five cases with results and analysis for each. Before any of the cases are discussed, the design parameters of the algorithm are chosen and motivation is given for the choices. The first case is a noise variance estimation problem with varying measurement noise levels on all satellites. The second introduces the offset detection problem by first showing a simulation of a no-fail condition. The third is an offset detection problem looking at instantaneously applied constant biases (step offset errors). The fourth study deals with errors that increase with time (ramp offset errors). The last case discusses the combined problem of offset error detection in a time-varying measurement noise variance environment. This last problem has long been considered the most difficult to overcome.

4.2 Design Parameters

The error detection algorithm inherently has many variables which, when changed, affect the detection performance of the algorithm. The variables are described and assigned values in this section for the simulations in this chapter.

4.2.1 Criteria

The goal of the tuning was to achieve no false alarms during nominal operation while minimizing the missed alarms during the application of step offsets, described in Section 4.5. This criterion was chosen to avoid nuisance alarm declarations during nominal operation. Clearly, any criteria could be employed in setting these parameters, as to achieve no missed alarms for some prespecified offset while allowing a minimal but non-zero false alarm rate. This might be useful if it were deemed critical to avoid a missed alarm on true offsets while accepting some false alarms, since the latter would simply entail switching from the current four satellites being used to a *different set of four.* In these simulations, when a satellite is declared failed, that satellite is removed and the algorithm continues using only three. For that reason, a criterion is used which yields no false alarms.

4.2.2 Estimation Window Length

The measurement noise variance estimator calculates the noise statistics over the past N_R sample times. When $\mathbf{R}$ is unchanging, a smaller N_R produces a poorer estimate, i.e., one with a larger the error variance in the estimate itself. If a very low estimate is formed, the chi-square variable can grow to a point where a failure is declared falsely. This is of concern if $\mathbf{R}$ undergoes a large step increase, and the estimate, $\mathbf{R}$, lags behind appreciably due to a large N_R choice. More acceptable estimates of a constant $\mathbf{R}$ can be achieved with $N_R = 50$ (Cases II-V). The confidence of noise strength estimates over only 10 sample periods is low, but this value provides extremely quick response to step jams (Case I).

4.2.3 Covariance Tuning Parameter

Instead of setting the filter measurement noise variance R_f to exactly $\hat{R}$, a constant R_o can be added, as shown in Section 3.6.5. This decreases the inherent lag of the response of the noise estimation algorithm and should be set to complement N_R. Together, these two parameters can reduce the false alarm rate caused by a poor noise strength estimate. For Cases II-V, a value of 7 ft^2 is added to each estimated noise variance. Case I uses a value of 0 ft^2.

4.2.4 Filter Reset Period

The length of time the detection filters propagate before being reset to the estimate of the navigation filter is designated as T. Each filter runs for T seconds before its state estimates, covariance information, and chi-square variable are reset. It is important to know that the chi-square variable must be reset to zero after a reset since it now assumed all measurement up to the current time have been good. T is set to 18 seconds for every case in this work. Comparisons between different values of T are left for future work.

4.2.5 Chi-square Summation Length

The chi-square variables are sums over a window of the past N_χ seconds. N_χ, the chi-square summation length, must be less than or equal to T due to the resetting of the chi-square variable. Otherwise, improper data is being incorporated into the hypothesis test. The chi-square summation

length can be thought of as a low-pass filter for the hypothesis test. N_χ is set to 18 samples to minimize the false alarm rate.

4.2.6 Failure Threshold

The threshold or limit, L, at which a failure is declared directly affects both the false alarm and missed alarm rates. By choosing a constraint of no false alarms, the threshold is set to 64. This value causes no false alarms during 10 runs, each 100 seconds long, under a no-fail condition. Much work could be done in the future to find the right combination of all these parameters for optimal detection capability in certain environments.

4.2.7 Summary

The final parameter values for Cases II through V are listed in Table 1:

Symbol	Description	Value
N_R	Noise variance estimation window length	50
R_o	Measurement covariance tuning parameter	7 ft^2
T	Filter reset period	18 sec
N_χ	Chi-square summation length	18
L	Failure threshold; chi-square limit	64

Table 1. **Error Detection Parameters used in Cases II - V**

4.3 Case I: Time-Varying Measurement Noise Covariance

The noise variance estimation problem is the simpler of the two types of failures to detect. Also, the satellites which experience varying measurement noise variance need not be removed from the solution. By adapting the filter models, the contribution of the measurement is only degraded. Statistics are calculated over the window length of N_R measurements, the only user-defined variable in the estimation portion of the algorithm. An experimental value of N_R=10 is used in this case. This value is shown to respond fairly quickly to step changes in true variance magnitude, as well as providing acceptable estimates. The algorithm is tasked to estimate the four diagonal elements of the **R** matrix, or the measurement noise variance of each individual satellite. These estimated diagonal elements will be designated by the scalars, R_i, where $i = 1, 2, 3, 4$. The off-diagonal

elements represent the cross-correlation of the white noise processes of two satellites. These terms are very small in reality and are assumed zero for this study. This yields an $\hat{\mathbf{R}}$ matrix as follows:

$$\hat{\mathbf{R}} = \begin{bmatrix} \hat{R}_1 & 0 & 0 & 0 \\ 0 & \hat{R}_2 & 0 & 0 \\ 0 & 0 & \hat{R}_3 & 0 \\ 0 & 0 & 0 & \hat{R}_4 \end{bmatrix} \tag{57}$$

Figure 4 graphs the square root of the diagonal element of $\mathbf{R}$, true standard deviation, as a solid line; the Monte Carlo (MC) mean of the estimate of that element as an 'x'; and the mean plus and minus the MC standard deviation of the estimate as a dotted line. The y-axes are the magnitude of the measurement noise standard deviation in feet. All results in this chapter are products of ten-run MC simulations over 100 second spans. The single case shown in this section is representative of all results yielded by this algorithm, i.e. no coupling between satellites and a response time directly related to the window length. It should be noted that interference power levels are correlated with white noise strength [18]. These simulations are intended to represent varying levels of in-band, or out-of-band, interference to the GPS signals, as well as failed satellites.

A variety of noise level sequences are shown to demonstrate the independent nature of the algorithm and its performance under different situations. This is simulated to represent real-world scenarios since low elevation satellites have significantly higher measurement noise variance and jammers might affect only limited SVs [7]. Satellite vehicle (SV) 1 undergoes a series of increasing steps to approximate an approach to a jamming site. Both SV2 and SV4 step up to a constant level and then back down. SV3 shows a very difficult estimation problem for this algorithm. The steps alternate at exactly the window length period. A faster estimator response to true changes could be achieved with a smaller N_R, but at the expense of larger error variance in that estimate. The parameters used for this case are N_R =10 and R_o=0 ft^2.

The navigation estimates of a KF with and without the aid of the noise variance estimate are shown in Figure 5. By comparing the figures, the benefit of the adaptive algorithm is evident. During periods of higher levels of measurement noise, the adaptive filter informs the user of this

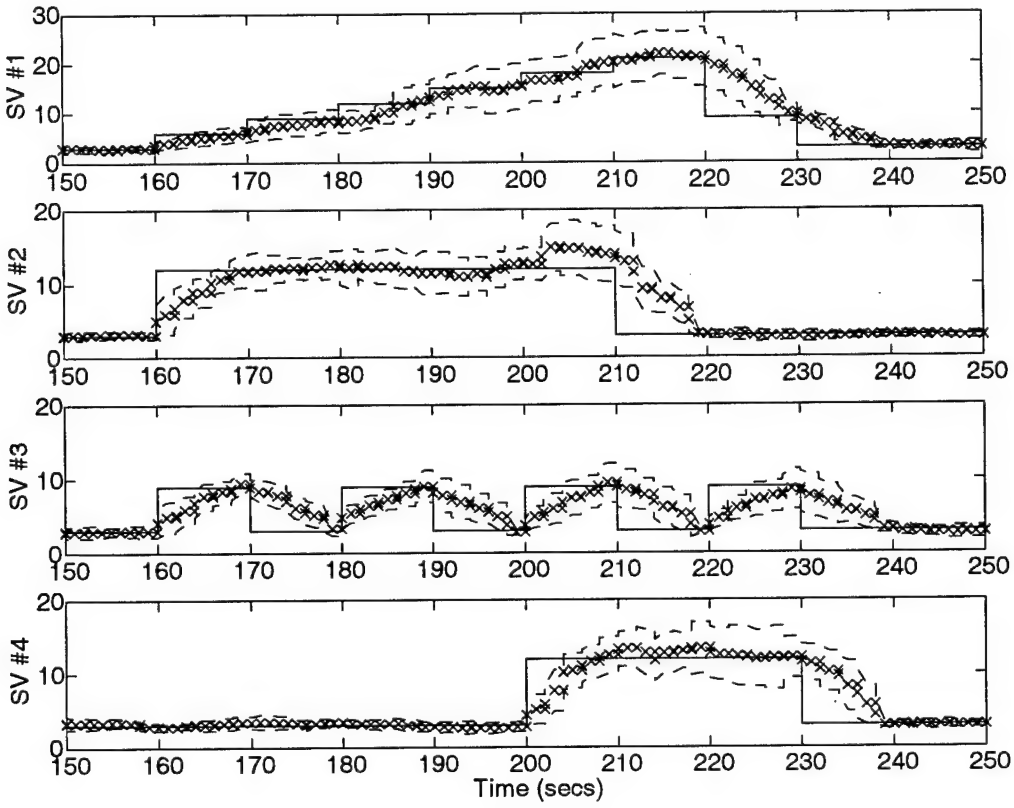

Figure 4. **Case I: True (solid line) vs. Estimated (mean 'x' $\pm$ 1σ 'dashed line') Measurement Noise Standard Deviation (ft)**

change by estimating the true error variance. The correction of the filter's internal model reduces the error committed by the filter. These figures show the ensemble mean over ten runs as a solid line; the mean plus or minus the standard deviation of the 10 runs is shown as dashed lines; and the filter-computed standard deviation is shown as solid lines centered about zero.

The non-adaptive filter underestimates the true standard deviation by a factor of three to four at times during the simulation. The adaptive filter clearly outperforms the non-adaptive filter, due to the accurate measurement noise estimation seen in Figure 4.

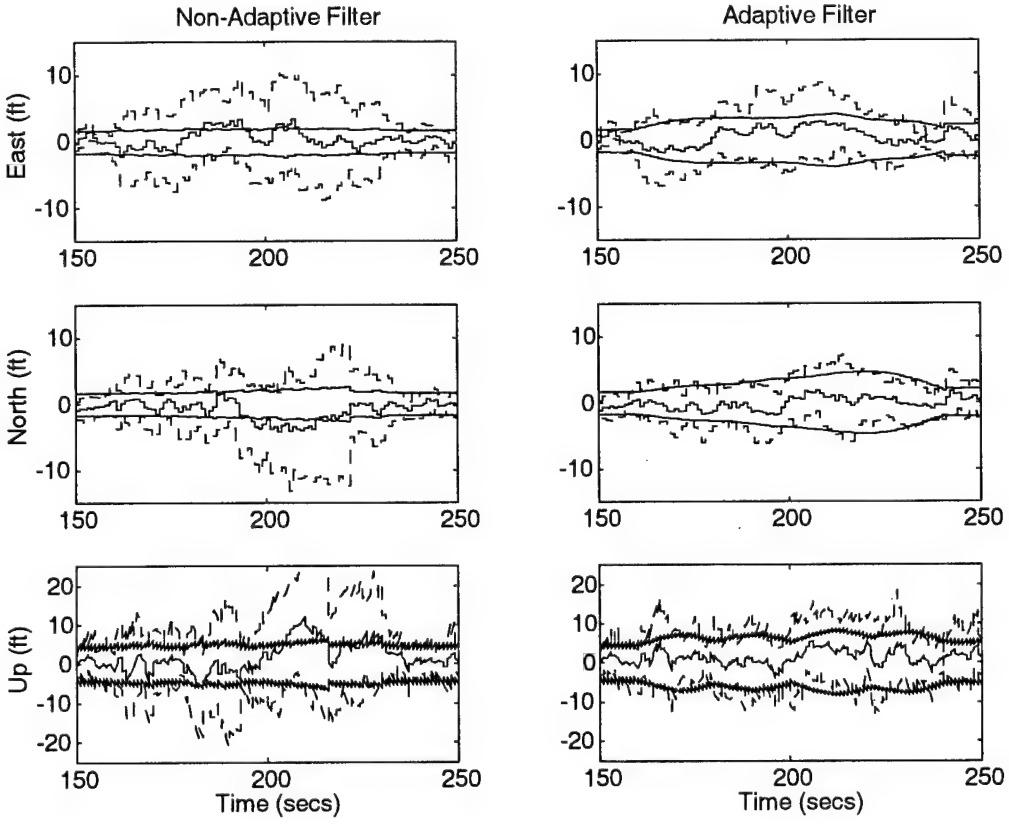

Figure 5. **Case I: Navigation Estimate Errors During Varying Noise Levels; plots show ± filter-computed standard deviation (solid), Monte Carlo mean (stairstep), and Monte Carlo mean ± standard deviation (dashed)**

4.4 Case II: Nominal Operation

This case describes the offset error detection process during a period of time when no errors are induced on the GPS signals. The output data is shown on a set of two figures and then discussed. The first, Figure 6, contains both residual and chi-square information of the satellite under test from the detection filters. Each plot on the left shows the residuals' MC mean (solid), MC mean $\pm$ standard deviation (dashed), and associated zero $\pm$ filter-computed residual standard deviation (solid), $\sqrt{A}$, from the detection filters. Note how closely the filter-computed $\sqrt{A}$ matches the residual mean $\pm 1\sigma$ curves. These plots refer to the satellite under test, which is the only satellite subjected to the simulated failures described in this chapter, i.e., *single* failures are induced on the only satellite hypothesized to fail. The information for the other satellites is not shown for any of the cases since it did not exhibit any characteristics dissimilar from those shown. The set of three plots on the right in Figure 6 show the chi-square variables from each detection filter for the satellite under test. Each MC run is shown as an individual curve. This is done to depict the variation in the runs visually as well as the absolute minimum and maximum values. The first two sets of three plots are adjacently displayed in Figure 6 due to the close relation of their data. The horizontal lines across the chi-square plots represent the threshold at which a failure is declared.

Figure 7 contains a set of three plots that portray the navigation estimate errors of the navigation filter. All three channels, East, North, and Up, are shown. The last plot contains the mean failure declaration information. When a satellite is declared failed, its associated 'figure of merit' (FOM) is changed from one to zero. The MC mean of a satellite's FOM will be one for all time in this case, but, in general, will vary between one and zero by increments of 0.1 since there were 10 MC runs performed. The mean FOM will only decrease since satellites are not brought back online. This number clearly shows the false alarm rate, deviation from one before the failure, and the missed alarm rate, deviation from zero after the failure. Note that, due to the threshold value chosen, there are *no false alarms* declared, and thus the mean FOM remains constant at unity during periods without failures.

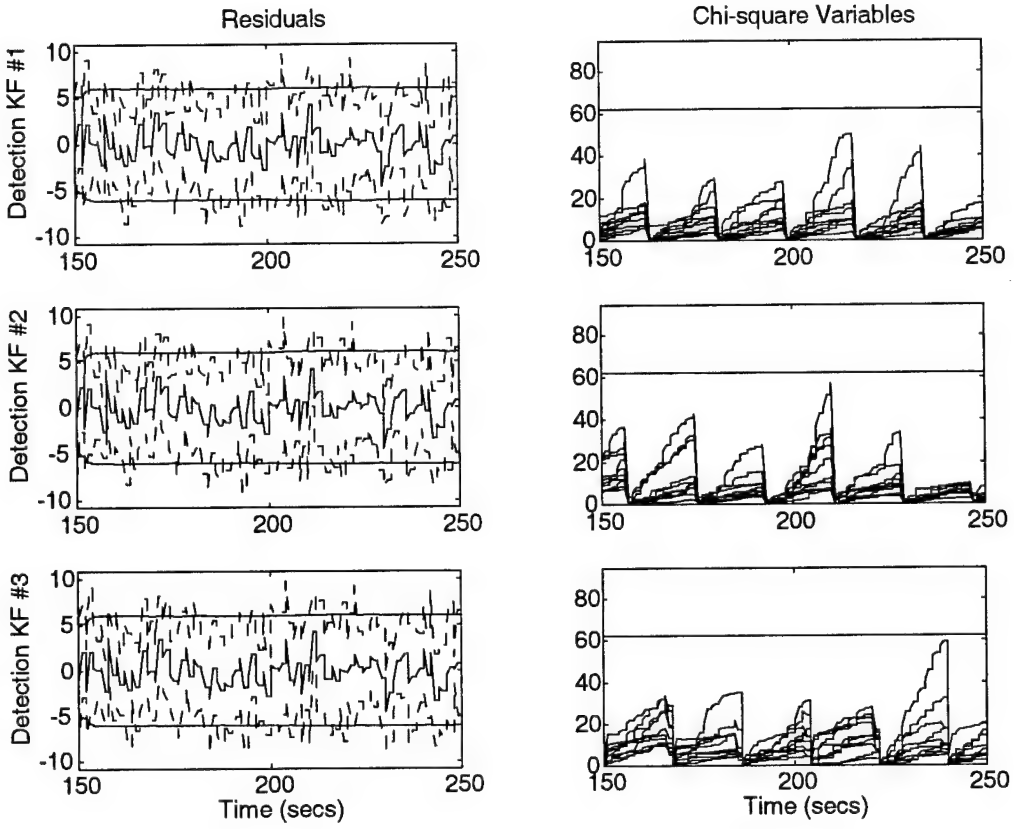

Figure 6. **Case II: Residual and Chi-square Information of the Satellite under Test during Nominal Operation**

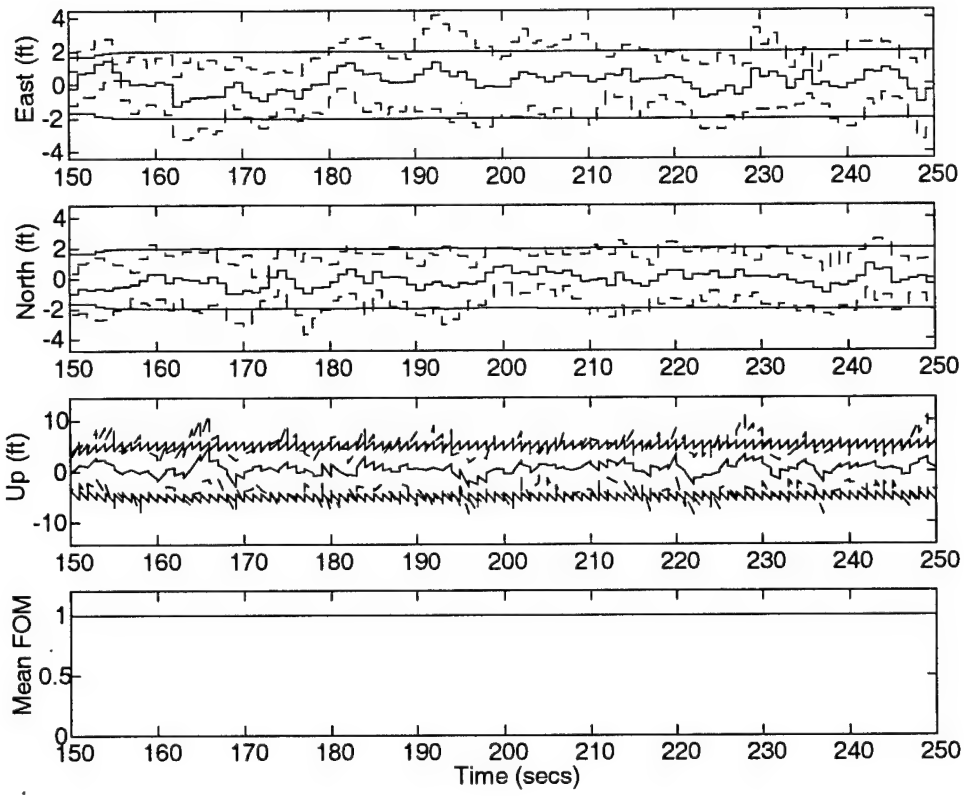

Figure 7. **Case II: Navigation Estimate Errors and Mean Figure of Merit Information during Nominal Operation**

4.5 Case III: Step Signal Offsets

The third case simulates the failure of a single GPS satellite at t=200 seconds by adding a step offset to the range measurement of one SV. This line of sight offset will enter all three navigation channels of the filter. Using the tuning values shown in Table 1, missed alarms do not occur for a step offset of 22 feet. The detection results for this case are shown in Figure 8.

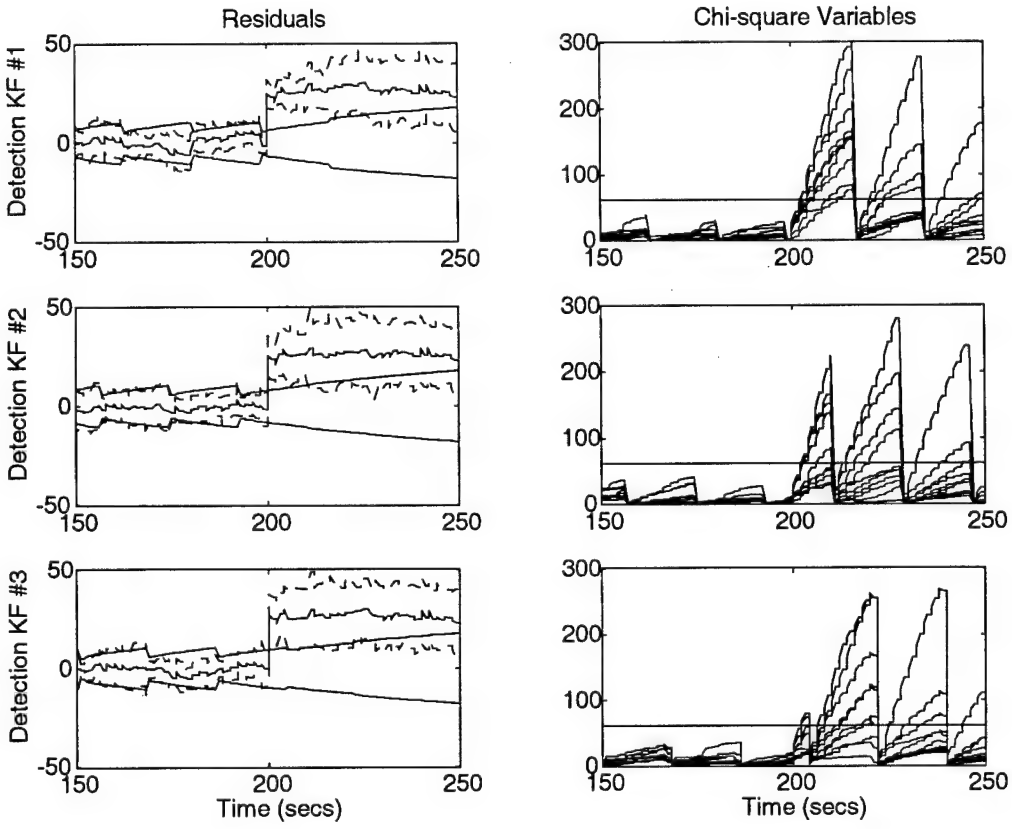

Figure 8. **Case III: Residual and Chi-square Information with 22' Step Offset Induced at t=200 sec.**

There is extensive information contained on the plots in Figure 8. The mean $\pm$ 1σ statistics of the residuals calculated by the detection filters for the single satellite under test are shown on the left hand side of the figure. The filter-computed $\pm$ standard deviation of residuals is also shown on these plots as smooth solid lines centered about zero. The associated chi-square variables are shown on the right of the figure. Note how conservatively the threshold is set, shown as the horizontal line

on the chi-square plots. This ultimately yields a zero false alarm rate for every study included in this thesis. The chi-square plots clearly show the reset timing of the three detection filters as well as the growing length summation property of the chi-square variable.

The residual information from the detection filters clearly indicates a failure at 200 seconds. The residuals of a single non-adaptive KF which is being updated every second by all four satellites, including the failed satellite, would not point out the failure with such clarity. A single filter experiencing an identical failure at t=200 seconds is simulated for comparison. The results are shown in Figure 9. Notice that the step is immediately detected in the residuals, but is soon lost as the filter states are corrupted.

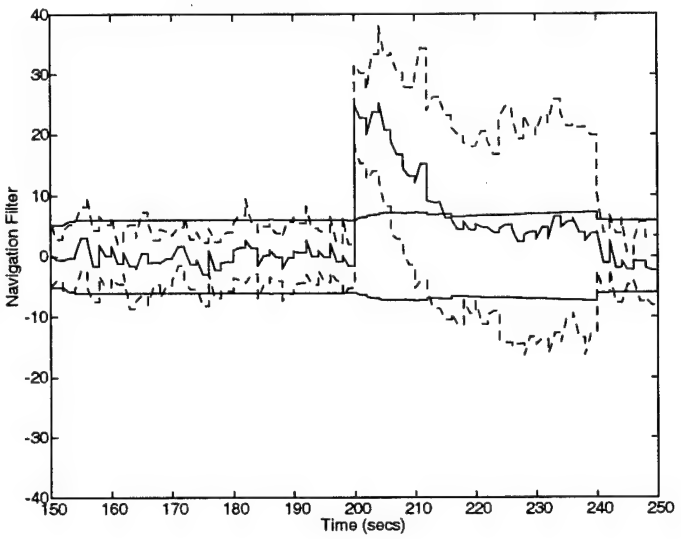

Figure 9. **Case III: Residuals from Navigation Filter with 22 ft. Step Offset Induced at t=200 sec.**

The navigation error plots, shown in Figure 10, demonstrate how the biased range measurement affects all the navigation estimates of the non-adaptive filter. The term 'adaptive' in the plot titles refers to the filter structure described in this thesis which contains multiple hypotheses to isolate failures. The 'non-adaptive' results are obtained from a single KF without any additional logic. Notice how the filter-computed standard deviation of the adaptive filter starts to grow after the failure occurs. This happens since the measurement from the failed satellite is not used in the

update of the navigation filter, i.e., the satellite which is declared failed is removed and only three satellites are used instead of four. This error growth does *not* continue unbounded. The error variance increase is strictly due to the loss of one satellite and it has almost reached steady state at the end of the simulation. In a real-world implementation of this algorithm, there would be *additional satellites available for testing and incorporation after a failure.* By replacing the GPS measurement after the failure with a measurement from a *fifth* satellite that was not in the original set of four, the error characteristics would return to the original non-failed values. Of course, these previous statements are only true if the algorithm detects the failure so that the measurement can be replaced.

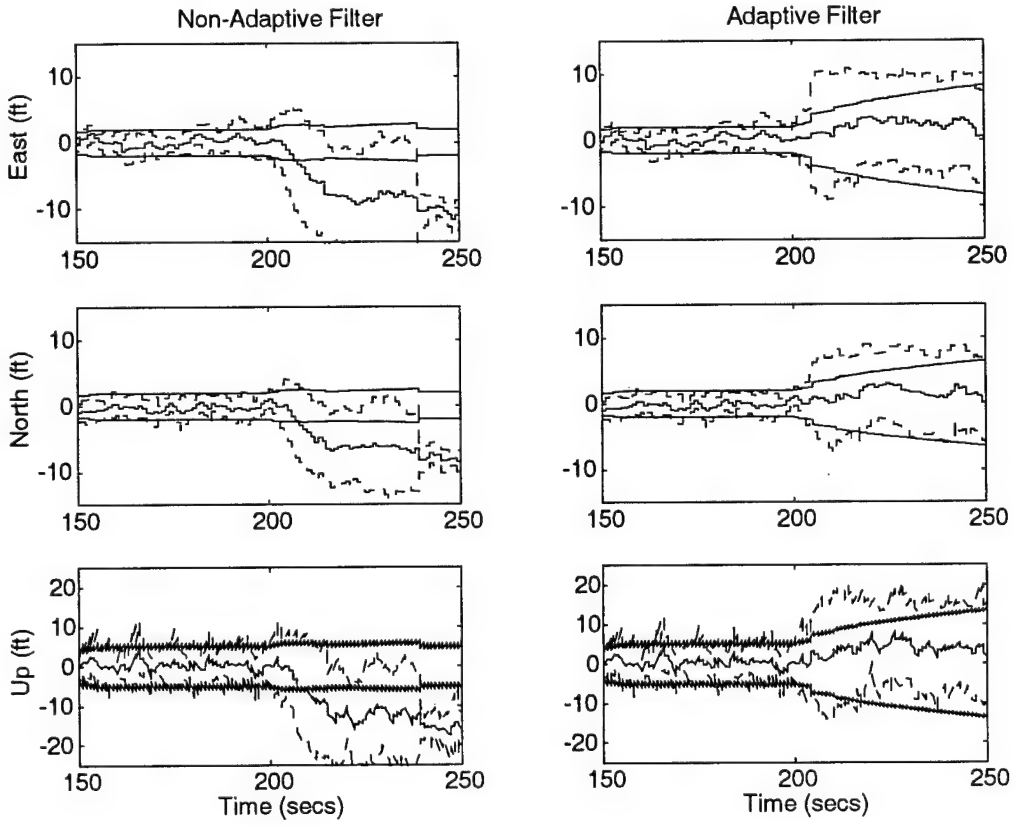

Figure 10. **Case III: Navigation Estimate Errors with 22' Step Offset Induced at t=200 sec.**

When various magnitudes of offset are applied to the GPS signal, the results shown in Figure 11 and summarized in Table 2 are attained. Even though the results shown for step offset detection meet with expectation, a better test for step offsets would have a shorter window to form the chi-square variable, i.e. a smaller N_χ. Also, setting a tighter threshold, L, than that listed in Table 1 would allow for lower missed alarm rates (and a zero missed alarm rate for less than the 22 ft. offset seen in Table 2), at the expense of a non-zero false alarm rate. The parameters used for all these cases are *optimized to detect slowly increasing, ramp-like offsets.* Also, realize this algorithm uses a single test to detect errors. Additional tests could be added to exploit other available data. This idea is expounded upon in Chapter 5.

Step Magnitude (ft)	Missed Alarm Rate	False Alarm Rate
22	0	0
20	0.1	0
15	0.2	0
10	0.7	0
5	0.8	0
0	N/A	0

Table 2. **Step Offset Detection Results Summary**

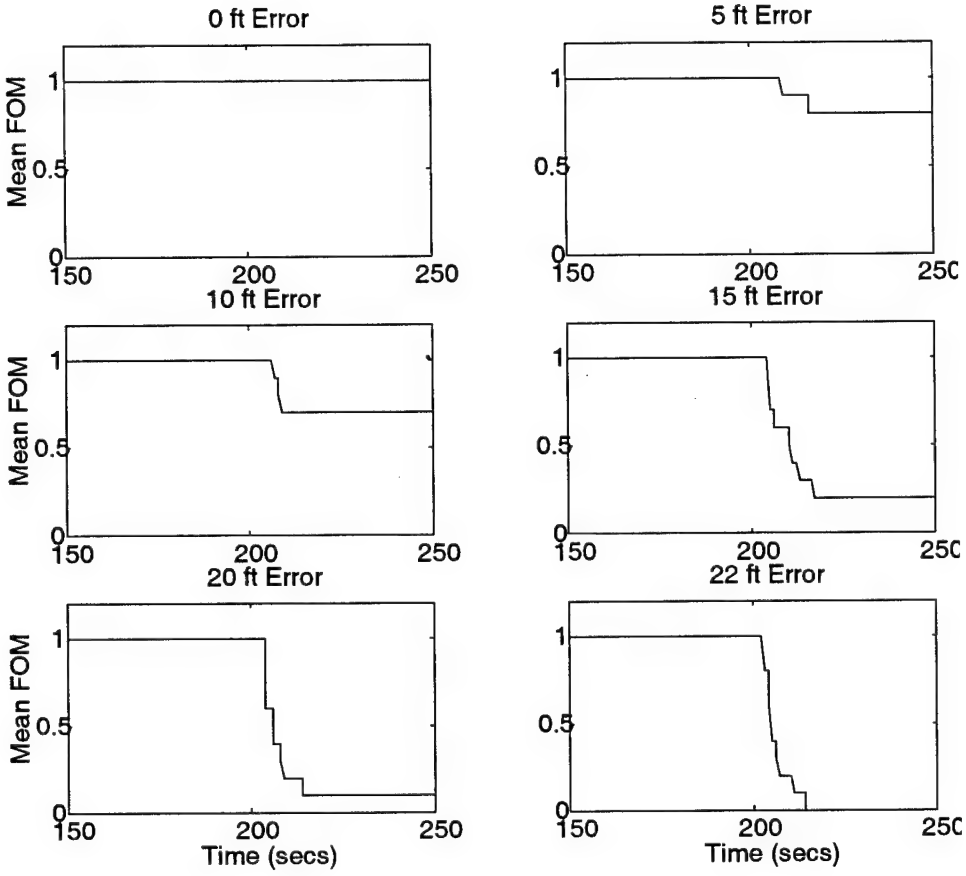

Figure 11. **Case III: Step Offset Detection Results**

4.6 Case IV: Ramp Signal Offsets

Case IV induces an offset error on a single satellite starting at t=200 seconds. This time the offset is an increasing ramp of constant slope. The offset runs for a total of fifty seconds, until the end of each simulation run. This means, for a ramp of 2.0 ft/sec, the offset is at a magnitude of 100 ft by the end of the run. These errors become much larger than the step offsets, but are generally more difficult to detect since there is no sudden change in residuals. This ramp type of error simulates the kind of errors which will be encountered in the real world: a satellite clock goes awry and begins to drift much faster than anticipated, a satellite begins to drift off its normal orbit, or a purposeful ramp is inserted by someone attempting to deny the user of accurate GPS position data.

Figure 12 demonstrates in practice how this algorithm is oriented to detect ramps. Since there is a slow, continual movement of residuals in one direction, the filter that gets updated once a second only sees one or two feet of error every update, which is not a large enough change to cause alarm. On the other hand, the filter which does not use a certain measurement for a longer period of time can tell that it is moving away from agreement with other devices, i.e. eighteen or thirty-six feet over eighteen seconds.

Again, as in Case III, the residuals of the detection filters are obviously showing the offset. A comparison to a single non-adaptive filter using all SVs to update its estimate in shown in Figure 13. The single filter corrupts its position and velocity states to agree with this erred measurement over time. The residual calculations shown in this plot are conditioned on the failed satellite as well as the three unfailed SVs. This is the main difference between a single filter application and the detection filters.

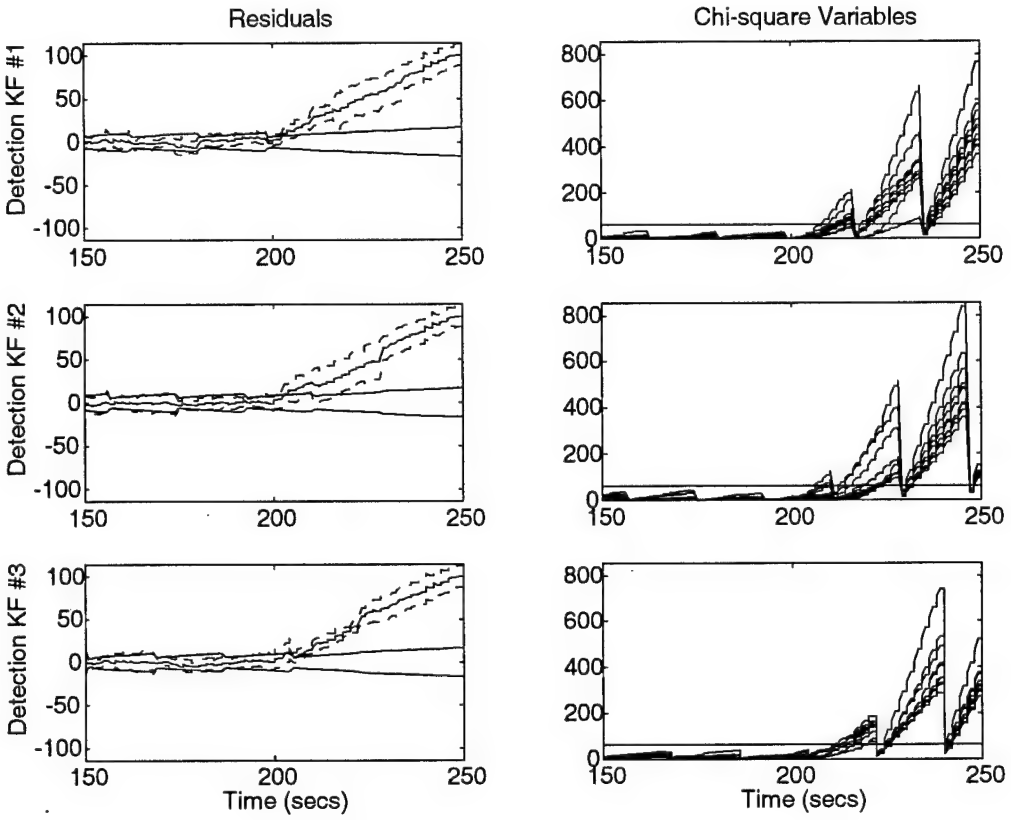

Figure 12. **Case IV: Residual and Chi-square Information of the Satellite under Test with 2 ft/sec Ramp Offset Induced at t=200 sec.**

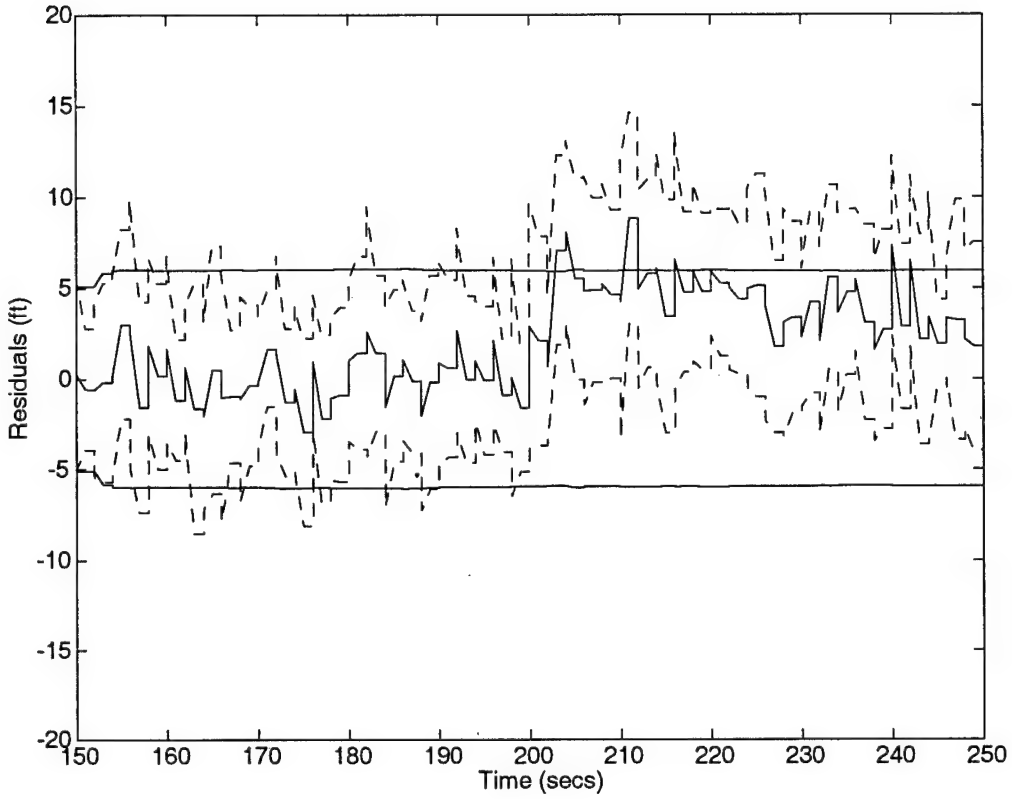

Figure 13. **Case IV: Residuals of Failed Satellite from Navigation Filter during Induced 2 ft/sec. Ramp Offset at t=200 sec.**

Figure 14 shows the navigation results for an offset ramp of 2.0 ft/sec. Remember, as in Case III, the error variance growth during the second half of the run for the adaptive filter is due to the declared loss of a satellite. With additional satellites available, this growth could be reversed back to the steady state statistics during the first half of the simulation. Notice that the adaptive filter uses the measurements from the bad satellite for a short time after the failure onset and therefore becomes corrupted. Then notice how after the failure is declared, the navigation filter is reset by the detection filters to an *uncorrupted* estimate. This is another highlight of the algorithm under evaluation: it *retrieves an uncorrupted estimate after failures are declared.* This study takes the most conservative approach and resets the navigation filter with the estimates of the detection filter which has been running the longest time without a reset. The past 18 seconds of measurements from the failed satellite are effectively removed from the estimate in the filter when a reset occurs.

The results of applying various offset magnitude rates to a single satellite's measurements is graphically displayed in Figure 15 and also listed in Table 3. The 0.5 ft/sec offset is equivalent to a drift rate of about 0.3 nmi/hr. Realize that the INS in effect is receiving the benefit of three satellites which bounds the maximum drift that occurs. Without this benefit, the detection filters could not be expected to come close to these results with an INS of similar quality, 0.42 nmi/hr for this simulation. For one run, the offset was not detected before a full filter reset period and corruption of the navigation estimate occurred. This is seen clearly in Figure 15 for the 1.5 ft/sec case. Table 3 reports the missed alarm rate (at 250 sec) for this offset ramp magnitude as zero, but note that the last declaration occurs at nearly 250 sec. When the failure was declared and the navigation filter reset, its estimate was improved significantly, but still left with a small offset. This occurs for any detection after 218 sec, since N_x is equal to 18 and the failure occurs at 200 sec.

Ramp Magnitude (ft/s)	Missed Alarm Rate	False Alarm Rate
2.0	0	0
1.5	0	0
1.0	0.3	0
0.5	0.7	0

Table 3. **Ramp Offset Detection Results Summary**

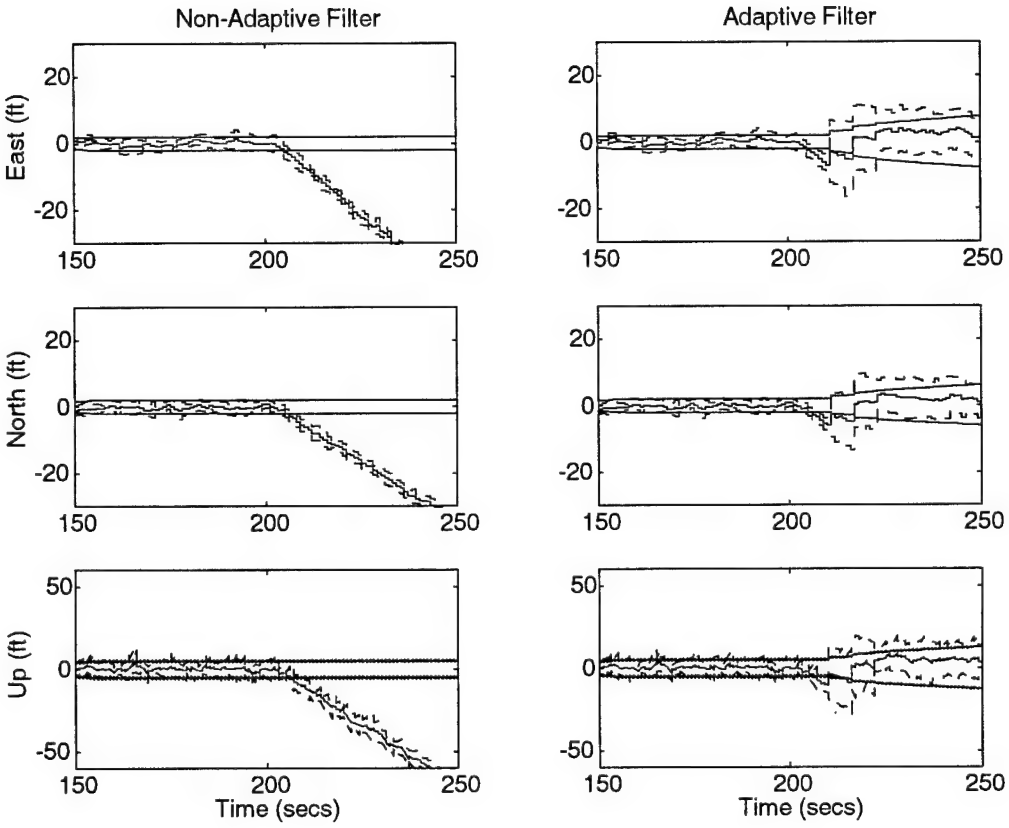

Figure 14. **Case IV: Navigation Estimate Errors during 2 ft/sec Offset Ramp Induced at t = 200 sec.**

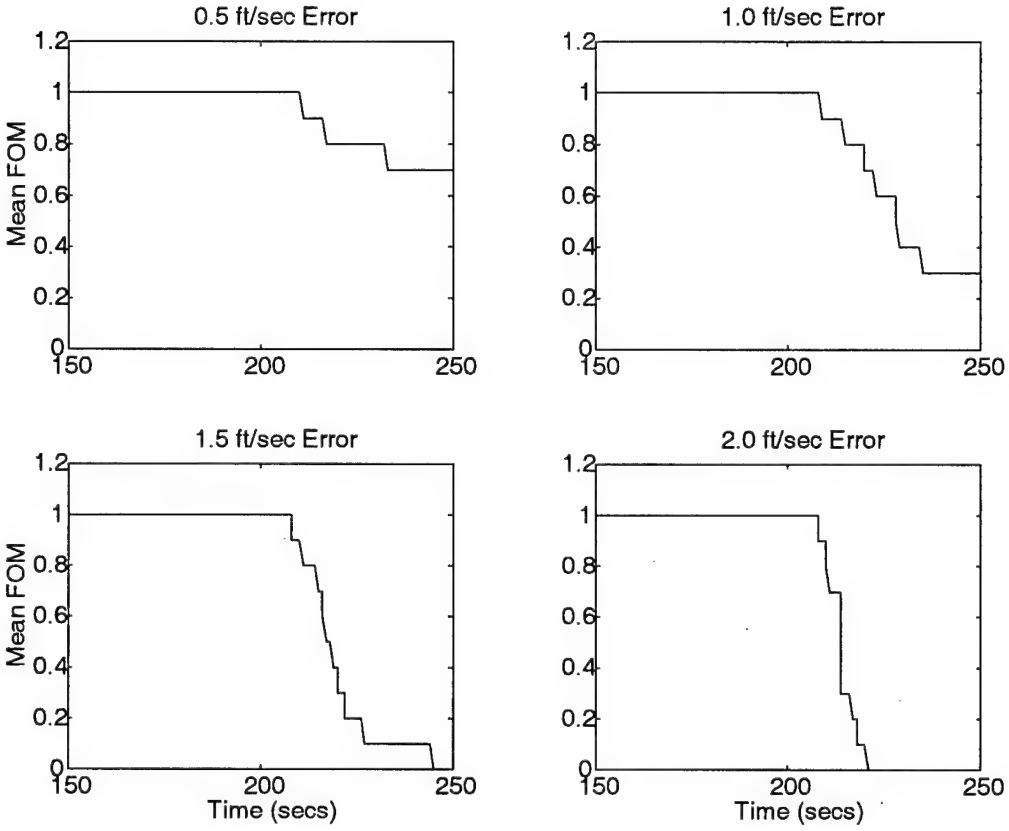

Figure 15. **Case IV: Ramp Offset Detection Results**

4.7 Case V: Ramp Signal Offsets during Increased Noise Variance

Case V induces the same ramp offsets as in four, but raises the difficulty of detection by increasing the measurement noise variance before the time of failure. The measurement noise is increased to a level of three times the nominal standard deviation for a period of fifty seconds before and thirty seconds after the onset of the ramp. This yields an $\hat{R}_i$ of 81 ft^2 for every satellite, compared the nominal value of 9 ft^2. Figure 16 shows the residual and chi-square information for the 2.0 ft/sec offset. These results are quite similar to those of Case IV. The algorithm seems to be able detect this error just as easily as it did with a much lower measurement noise variance.

The first thing to notice about the navigation plots in Figure 17 is that the scales are different between the non-adaptive and adaptive filter plots. This is done to show the significant increase in navigation error variance throughout the simulation for the non-adaptive filter. The abysmal performance of this filter is due to the fact of it having two incorrect assumptions: a low measurement noise variance and the presumed non-existence of signal offset in the 'failed' satellite measurement. The filter-computed standard deviation of the adaptive filter may seem to be having more difficulty matching truth than in previous case, but this is an artifact of the detection taking a little longer for some runs in the simulation. The runs in which a failure is not declared until 230 seconds do introduce some corruption into the navigation filter when the reset takes place. When the statistics of the ensemble are formed, the slow detections, and the associated errors, are merged with the quick detections, with no errors. The issue of corrupted resets brings to light an important point: there is a trade-off between reset estimate quality and the length of time corrupted measurements can go undetected and still be successfully removed when a failure is declared. The detection filters are reset to decrease their error covariance, which in turn increases the quality of a reset, if one occurs; but this then limits how long a period a reset can 'rewind' the navigation filter. It would be possible to run an additional bank of filters with a much longer reset period. This configuration maintains a high-quality estimate which is recalled in the event of an easily detected error to reset

the navigation filter. In addition, an older estimate is maintained in the event of a failure which stays undetected for a long period of time. These ideas are explored further in Chapter 5.

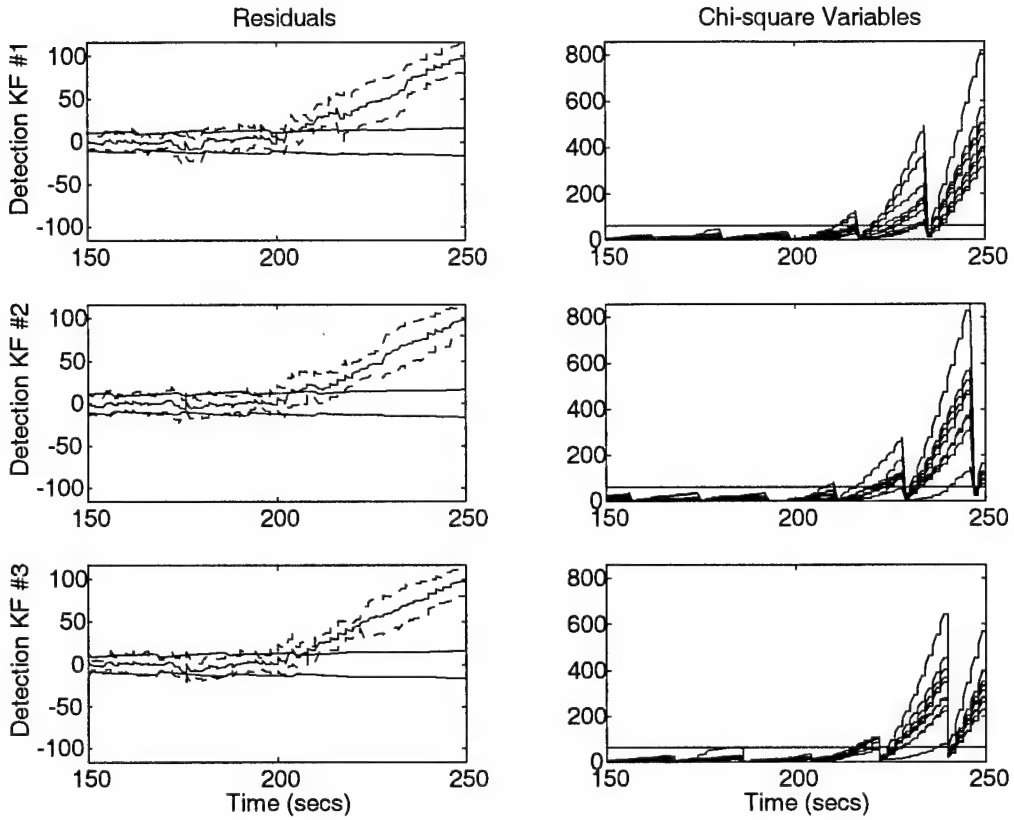

Figure 16. **Case V: Residual and Chi-square Information**

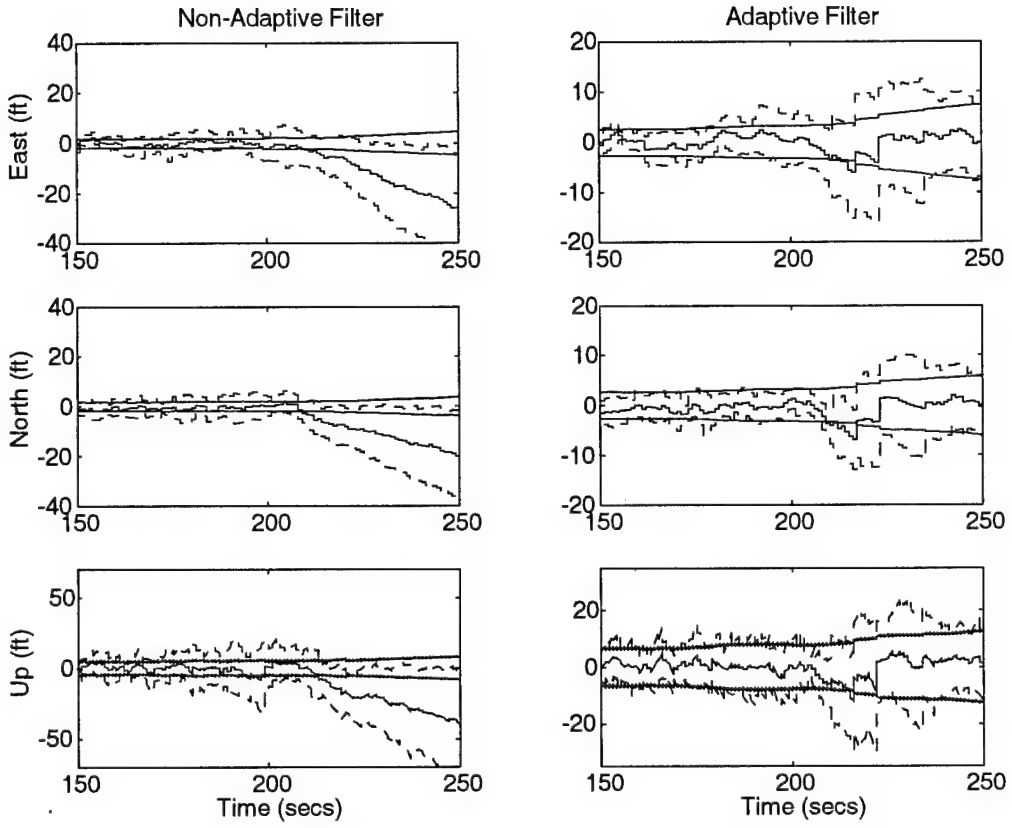

Figure 17. **Case V: Navigation Estimate Errors for 2'/sec Offset Ramp during Increased Measurement Noise**

Figure 18 shows, and Table 4 lists, the results for the application of different offset rates. It should be evident, when compared to Figure 15 in Case IV, that the increase in measurement noise variance does effect the detection threshold of the algorithm. This result should be expected, although it is worth noting that, in the performance above, the detection threshold seems relatively unaffected. This is a very desirable quality and suggests that even under worse measurement noise conditions some offset error detection capability remains.

Ramp Magnitude (ft/s)	Missed Alarm Rate	False Alarm Rate
2.0	0	0
1.5	0	0
1.0	0.7	0
0.5	0.9	0

Table 4. **Ramp Offset Detection Results during Increased Measurement Noise**

4.8 Summary

This chapter describes the capabilities and characteristics of the parallel KF structure described in Chapter 3. The measurement noise variance estimator and the offset detection algorithm were tested through simulations. The first four cases separately analyze their performance and the last combined both measurement corruptions to demonstrate the joint ability of the entire algorithm. The simulations described here warrant further investigation of this parallel filter structure to exploit its full capabilities. The recommendations for further work are included in Chapter 5.

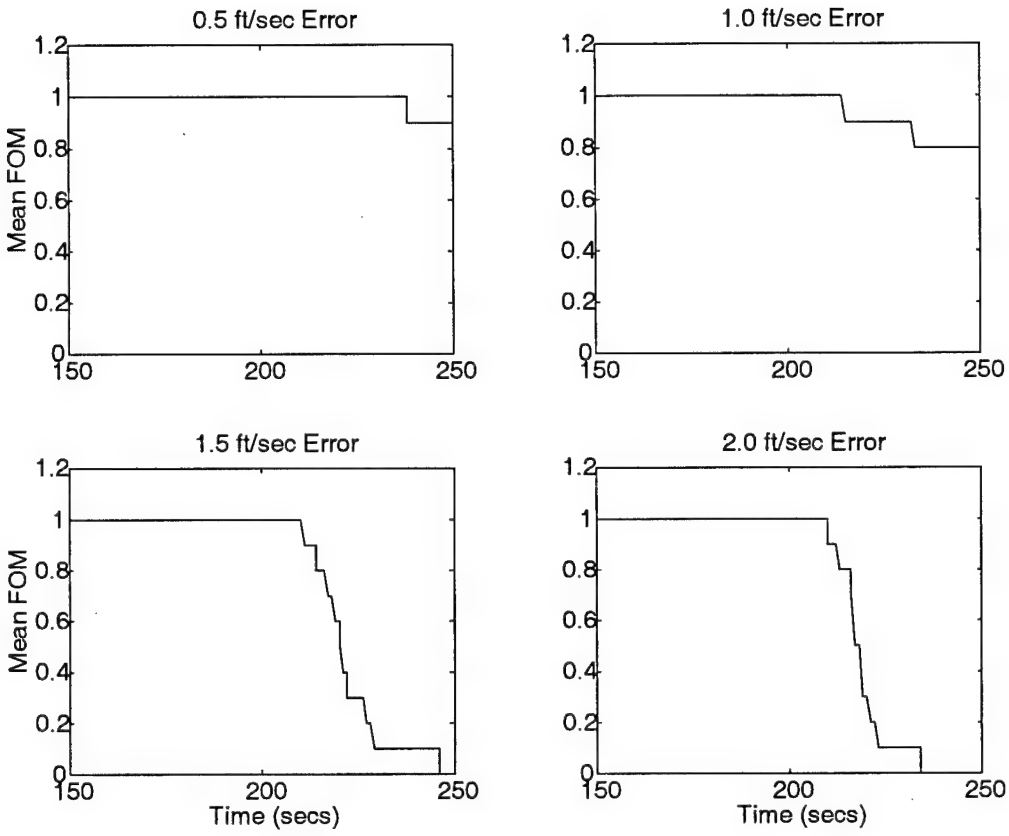

Figure 18. **Case V: Detection Results for Offset Ramp during Increased Measurement Noise**

Chapter 5 - Conclusions and Recommendations

5.1 Conclusions

The results presented in this thesis sufficiently demonstrate the viability of the offset detection algorithm with decoupled variance estimator. It should be noted that any reliable detection is an improvement over the standard system which has *no* detection ability. Chapter 4 includes performance analyses which list the degree of success for various offset magnitudes and magnitude rates. The simulations are only performed for one set of parameters; therefore, these results are dependent on the parameter choices made at the beginning of Chapter 4. Modifications to these values would certainly affect the results, possibly in a positive manner. More importantly, only a single test statistic, i.e., a summed chi-square variable relating to scalar residuals, is used to declare failures. This choice alone has a large impact on the type of performance the algorithm will yield. In fact, the test statistic used in this research was designed specifically to detect ramp offsets.[8] The real intent of the research is to demonstrate that *this technique can adeptly detect the more difficult of the two failure types*, viz. ramp offsets. Chapter 4 demonstrates the ability to detect both step and ramp offsets with a single test statistic and to remove their corruptive effects from the navigation system outputs.

The detection algorithm is based upon the parallel detection filters which compute scalar residuals for measurements that have not been used recently to update the filters. This technique provides objectively formed residuals for use in hypothesis testing. The detection filters *do not corrupt their states* due to failures in the satellite under test; therefore, the residual calculations continue to identify the failure correctly for long periods of time. Summations of even small deviations can be detected over time. Normally, without the protection provided by this algorithm, a bad satellite will corrupt the state estimates in a conventional filter within a few time samples and the filter will never have another chance to detect the failure. Even if the failure is detected, the single filter cannot remove the effect of the bad measurements that were incorporated. Similar summations of

[8] A smaller value for N_χ should be chosen to improve the step offset detection.

test statistics from this conventional type of filter over time will not provide as clear an indicator. This technique not only provides a method to detect slowly increasing offsets, but also a method to *correct* the navigation estimate after the offset is detected and *remove* the effect of the corruption from the navigation system output.

This research is a *proof-of-concept* which is intended to demonstrate the characteristics of an algorithm. The summarized performance of this algorithm is only for one choice of parameters and should not be taken as a global measure of its capability. The characteristics described in the previous paragraph clearly indicate, all things equal, this parallel-filter algorithm will outperform any other residual monitoring of a single filter. However, because the satellite-under-test does not update the detection filter, an uncorrupted comparison between a single measurement and the best available estimate is obtained. As a result, the best available estimate in the detection filter does not gain the precision offered by including that single measurement.

This last statement exposes a limiting property of the algorithm. Detection filters, as currently formulated, do not remain uncorrupted in the event of simultaneous dual failures. Two satellites must be left out of the update process for this to occur. Likewise, to attain protection from n simultaneous failures, updates from n satellites must be withheld. Failures are not considered simultaneous as long as individual failures are detected before the onset of another.

The measurement noise variance estimation algorithm aids the detection process in these simulations, but it could also be implemented by itself. The redundant measurement differencing (RMD) technique provides direct observability of the uncorrelated measurement noise. This facilitates the decoupling of the parameter estimation from the offset detection problem. Measurement offsets do *not* effect the estimation of the measurement noise variance, since such offsets would be cancelled out in the RMD process. In fact, the results in Chapter 4 reveal the ability to estimate the noise variance and simultaneously detect and remove the impact of ramp offsets in the measurement signal, thereby successfully addressing what is often considered to be the most difficult noise strength variation and measurement offset scenario available. As the algorithm stands now, the sliding win-

dow length completely determines the performance of the estimation process. Many modifications could be used to enhance the characteristics of this estimation process. Several are suggested in the next section. The success of the combined algorithm, even in a limited form, demonstrates capabilities which should be explored further. The purpose of the rest of the chapter is to suggest possible avenues of pursuit for this continued research.

5.2 Recommendations

5.2.1 Measurement Noise Variance Estimator

The estimator in this study is implemented in the simplest form possible. The more important aspect of the work is the redundant measurement differencing (RMD) technique which leads to these simple empirical variance calculations. These estimates tend to lag a time-varying **R**. This lagging causes false alarms when **R** increases or missed alarms when **R** decreases. The estimator can interpret the individual differenced noise samples in any desired manner. This thesis presents a simple sliding window technique whose estimates lag changes by exactly the window length. The first suggestion to is to weight the data according to age, i.e., weight the most recent differences more heavily than older measurements. This would enhance the response time of the estimate while maintaining the error variance performance. Also, computing the R_o of Section 3.6.6 adaptively could enhance the response time without lowering the fidelity of the estimation.

A real-world hardware test should be carried out to demonstrate the capability of the algorithm under true conditions. Some receivers already have the ability to set a channel to search for a specific satellite. Redundant measurements could be differenced and the measurement noise variance estimated. This test could be enhanced by introducing interference sources in the area around the receiver to induce increased noise variances. The validation of this algorithm with real-world tests is a critical step in the development of an operational system.

5.2.2 Offset Detection

By limiting the number of available satellites to four, the performance of the navigation filter is severely limited after a single failure declaration which reduces the number of satellites used from

four to three. The increase in standard deviation bounds is completely due to the reduction of the number of satellites used to update the filter. This reduction in performance forced a very conservative design constraint, viz. no false alarms. With the reduction to fewer satellites than four, false alarms would cause a degradation in navigation performance which would eventually affect the detection ability of the algorithm. With an increase of satellites in view, the navigation 1σ error bounds can be returned to the original level after a failure declaration by replacing the failed satellite with another available satellite not currently updating the filter. This allows a design approach which enables the false alarm rate to become non-zero in order to keep the *missed* alarm rate minimized. In final application, the missed alarm rate is more crucial than the false alarm rate, especially when additional satellites are available to regain an acceptable level of navigation performance after a false alarm.

More available satellites allow multiple failures to be simulated during a single run and, as satellites return to nominal operation, some criterion can be chosen to re-certify those previously failed satellites for use in the solution at a later time if required. However, to perform these simulations properly, the construction of a minimum of fifteen detection filters is required: twelve to monitor the four used in the solution and an additional three to monitor a 'spare' to include in the effect of a declared failure. Notice, an algorithm based on fifteen detection filters cannot experience two *simultaneous* satellite failures without a degradation in navigation performance. Nevertheless, two *sequential* failures, separated by a period of time sufficient to validate a different 'spare' before the second failure occurs, could be detected without any loss in performance. The increase of available satellites is only properly handled by creating a bank of three detection filters for every satellite that could be incorporated into the navigation filter, and assuming simultaneous multiple failures would not be a hypothesis of sufficient concern to warrant additional detection filters devoted to such an occurrence.

The addition of navigation aids to the simulation would allow an investigation of the detection ability of a more precise system. Pseudolites, a radar altimeter, or ground transponders could

be added to the model to characterize a landing profile or a flight test at a reference range. A better INS could be simulated as well. These system improvements would undoubtedly improve the detection capability of the algorithm. If a true performance analysis is desired, more realistic (higher order) INS and GPS models must be used as 'truth models' for simulating the real-world environment, rather than the simpler models used in this proof-of-concept study. This is the next step in validation through simulation.

A possible method to decrease the detectable magnitude of a ramp offset is to run additional filters in parallel which have much longer filter reset periods. Again, this requires additional filters, but improves the capability of the algorithm. As shown in Chapter 3, the longer the filter period is, the lower the detectable offset rate will be. Running two banks of filters for each satellite may be feasible, but might be equivalent to running multiple test statistics as described in the next paragraph.

A simple yet extremely valuable modification is the addition of test statistics to detect different types of failures. In Chapter 4, the results clearly show the summed chi-square test is tailored to detect ramp offsets, not steps. A new test statistic specifically intended to detect step offsets could be added to the algorithm quite easily. This test statistic could be the current chi-square variable with a smaller N_χ. This improves the performance when detecting step offsets and adds only a small amount of computational effort.

The number of filters seems to be ever increasing with every improvement proposed. The criteria for any good engineering design should include simplicity. In this case, a solution should use the minimum number of filters necessary to achieve desired results. One aspect of the detection algorithm, as it stands now, that creates additional complexity is the interleaved filters with staggered reset times. With the addition of more available satellites, it may be possible to make one change that reduces the complexity significantly and might improve the performance. This modification is to update each detection filter with four measurements while testing a fifth. If five satellites are considered in view, five different sets of four are possible that would each require *one* filter.

Only one is required since the precision of the estimate is not significantly[9] lower than that of the navigation filter. This allows the filters to run with no resets unless a failure is declared. Even with six available satellites, there are only 15 combinations of four. This is less than the 18 required by the current technique of using three per satellite and provides a better estimate with which to perform test statistics. More than six available satellites quickly increases the possible combinations of four significantly, but six available provides this new structure the capability to fail up to two satellites simultaneously without any degradation in navigation. Filters updated with fewer than four satellites could be included in this new structure as well to provide protection against the possibility of having fewer than four good satellites available at any one time. With newer GPS receivers having additional channels available, more than four satellites at a time can be used to provide better nominal navigation performance *and* better offset detection ability.

These recommendations are intended to give the reader insight into this detection problem. There is not a single best solution that will work for all cases. The filter structure and algorithm logic presented in this work are far from optimal, yet are meant to stimulate further research ideas in this area. In summary, the enhanced capabilities of future detection algorithms will be derived from multiple test statistics performed on the output of parallel filters.

[9]Some differences would occur due to geometry effects.

APPENDIX A - Simulation Models

A.1 Dynamics Matrix

$$\mathbf{F}_{(13x13)} = \begin{bmatrix} \mathbf{F}_{11} & \mathbf{0} & \mathbf{F}_{13} & \mathbf{0} & \mathbf{0} \\ \mathbf{F}_{21} & \mathbf{F}_{22} & \mathbf{F}_{23} & \mathbf{0} & \mathbf{0} \\ \mathbf{F}_{31} & \mathbf{F}_{32} & \mathbf{F}_{33} & \mathbf{F}_{34} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{F}_{43} & \mathbf{F}_{44} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{F}_{55} \end{bmatrix}$$

$$\mathbf{F}_{11} = \begin{bmatrix} 0 & 0 & -\rho_y \\ 0 & 0 & \rho_x \\ \rho_y & -\rho_x & 0 \end{bmatrix} \quad \mathbf{F}_{13} = \begin{bmatrix} 0 & -C_{RY} & 0 \\ C_{RX} & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}$$

$$\mathbf{F}_{21} = \begin{bmatrix} 0 & -\Omega_z & \Omega_y \\ \Omega_z & 0 & -\Omega_x \\ -\Omega_y & \Omega_x & 0 \end{bmatrix} \quad \mathbf{F}_{22} = \begin{bmatrix} 0 & \omega_{in_z} & -\omega_{in_y} \\ -\omega_{in_z} & 0 & \omega_{in_x} \\ \omega_{in_y} & -\omega_{in_x} & 0 \end{bmatrix}$$

$$\mathbf{F}_{23} = \begin{bmatrix} 0 & -C_{RY} & 0 \\ C_{RX} & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} \quad \mathbf{F}_{32} = \begin{bmatrix} 0 & -A_z & A_y \\ A_z & 0 & -A_x \\ -A_y & A_x & 0 \end{bmatrix}$$

$$\mathbf{F}_{31} = \begin{bmatrix} -2V_y\Omega_y - 2V_z\Omega_z & 2V_y\Omega_x & 2V_z\Omega_x \\ 2V_x\Omega_y & -2V_x\Omega_x - 2V_z\Omega_z & 2V_z\Omega_y \\ 2V_x\Omega_z & 2V_y\Omega_z & -2V_x\Omega_x - 2V_y\Omega_y \end{bmatrix}$$

$$\mathbf{F}_{33} = \begin{bmatrix} -V_zC_{RX} & 2\Omega_z & -\rho_y - 2\Omega_y \\ -2\Omega_z & -V_zC_{RY} & \rho_x + 2\Omega_x \\ \rho_y + 2\Omega_y + V_xC_{RX} & -\rho_x - 2\Omega_x + V_yC_{RY} & 0 \end{bmatrix}$$

$$\mathbf{F}_{34} = \begin{bmatrix} 0 & 0 \\ 0 & 0 \\ 2g_o/a & 0 \end{bmatrix} \quad \mathbf{F}_{43} = \begin{bmatrix} 0 & 0 & 1 \\ 0 & 0 & 0 \end{bmatrix}$$

$$\mathbf{F}_{44} = \begin{bmatrix} 0 & k_2 \\ k_1 & -\beta_{\delta h_c} \end{bmatrix} \quad \mathbf{F}_{55} = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}$$

$\rho_{x,y}$	Components of angular rate; nav to ECEF frame
$C_{RX,RY}$	Components of Earth spheroid inverse radii of curvature
$\Omega_{x,y,z}$	Components of angular rate; ECEF to inertial frame
$\omega_{in_{x,y,z}}$	Components of angular rate; nav to inertial frame
$A_{x,y,z}$	Components of specific force in the sensor frame
$V_{x,y,z}$	Components of vehicle velocity in ECEF frame
g_o	Equatorial gravity magnitude (32.08744 ft/sec^2)
a	Equatorial radius of the Earth (20,926,470 ft)
$k_{1,2}$	Vertical channel barometric altimeter gains (0, 0.03)
$\beta_{\delta h_c}$	Barometer inverse correlation time (1.6667e-3 sec^{-1})

A.2 Dynamics Driving Noise Strength Matrix

$$\mathbf{Q}_{(13x13)} = \begin{bmatrix} Q_{xy}\mathbf{I}_{(2x2)} & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & Q_z & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & Q_g\mathbf{I}_{(3x3)} & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & Q_a\mathbf{I}_{(3x3)} & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & Q_h & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 2\beta_{\delta h_c}\sigma^2_{\delta h_c} & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & Q_{\delta t} & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & Q_{\delta i} \end{bmatrix}$$

Q_{xy} PSD value of computer frame white noise (4e-15 deg^2/sec^3)
Q_z PSD value of computer frame white noise (1e-8 deg^2/sec^3)
Q_g PSD value of gyro drift rate white noise (6.25e-10 deg^2/sec^3)
Q_a PSD value of accelerometer white noise (1.037e-7 ft^2/sec^3)
Q_h PSD value of altitude error (15 ft^2/sec^3)
$Q_{\delta t}$ PSD value of GPS user clock phase (10 ft^2/sec^3)
$Q_{\delta i}$ PSD value of GPS user clock frequency (5e-12 ft^2/sec^5)
$\sigma^2_{\delta h_c}$ Variance of barometric altimeter correlated noise (1e4 ft^2)

APPENDIX B - GPS Error Models

B.1 Overview

The Global Positioning System (GPS) is a space-based radio navigation system. It is funded and controlled by the US Department of Defense (DOD). Nominal operation is with four satellites on each of six orbital planes for a total constellation size of 24. The satellites make two complete orbits in a sidereal day, about 23 hours 56 minutes. This configuration gives most points on Earth visibility to at least 5 satellites at any given time. Each satellite transmits a navigation message modulated on two frequencies and on two codes unique to each satellite. The two codes are the Coarse/Acquisition (C/A) and the Precise (P). C/A code's navigation message has intentional errors induced on it to degrade performance. This error source is named Selective Availability (SA). The Department of Defense employs SA to restrict access to high precision navigation data. Only certain receivers have the ability to use the P-code, as it is mainly used for military applications. The two frequencies are both in the L-Band and are named L1 and L2. L1 is the general use frequency, while L2 is available only to users with special receivers. The navigation message contains the information needed to calculate the range and range-rate to the satellite.

GPS can be used as a stand-alone navigation tool, but is often combined with an inertial navigation system (INS) for airborne applications. This fusion is performed mathematically with a Kalman filter (KF). The GPS ranges are used as measurements to estimate the errors in the INS. This combines the high fidelity of the INS in dynamic situations and the bounded error characteristics of GPS. Of course its performance is limited by how well the mathematical model in the KF matches reality. The task of modelling is always challenging and can become gruelling. Mother Nature is seldom well described by a simple differential equation and even less often 'ideal'.

Because the GPS signals propagate through the atmosphere and originate from multiple sources, there are many deviations from the simplest of theoretical models, viz. the range equals the speed of light multiplied by the time of transit. Because the range calculated by a GPS receiver always contains errors, it is called a pseudorange (PR) and it will denoted as ρ in this document. Differen-

tial GPS (DGPS) is a technique used to reduce some of the errors in the PR. This technique uses the difference between the range from a satellite to a user and the range from that same satellite to a known location as the measurement. Further mathematical detail is contained in this appendix. The following sections describe a simulation model used for the GPS signal.

B.2 Pseudorange Model

The GPS receiver calculates the range between the satellite, at the time of transmission, and the user, at the time of reception. This true range is given in vector form in Equation (58):

$$R_{true} = |\mathbf{r}_{sv}(t_s) - \mathbf{r}_u(t_r)| \tag{58}$$

The time arguments will be assumed throughout the rest of this work. The scalar calculation of range, shown in Equation (59), is performed in a rectangular frame, usually the Earth Centered-Earth Fixed (ECEF):

$$R_{true} = \sqrt{(x_{sv} - x_u)^2 + (y_{sv} - y_u)^2 + (z_{sv} - z_u)^2} \tag{59}$$

A spherical reference is desired for the purpose of navigation, viz. latitude or ϕ, longitude or λ, and altitude or h. This transformation, shown in Equation (60), assumes a WGS-84 Earth and any errors induced by this transformation are ignored in this study.

$$\begin{bmatrix} x \\ y \\ z \end{bmatrix}^{ecef} = \begin{bmatrix} (R_N + h)\cos\phi\sin\lambda \\ (R_N(1 - e_e^2) + h)\sin\phi \\ (R_N + h)\cos\phi\cos\lambda \end{bmatrix}^{ecef} \tag{60}$$

x_u, y_u, z_u	Cartesian coordinates of user
x_{sv}, y_{sv}, z_{sv}	Cartesian coordinates of a satellite vehicle
x, y, z	Cartesian ECEF triplet
e_e	Eccentricity of the Earth
a_E	Equatorial radius of the Earth
R_N	$a_E/\sqrt{1 - e_e^2 \sin^2\phi}$

The range calculated by the GPS receiver is by no means equal to the true range. It has various errors associated with it due to asynchronous clocks, atmospheric disturbances, electromagnetic effects, and the imperfection of real-world electronics. Equation (61) is the pseudorange model equation. It is extremely important to note that *all the errors have units of distance, even the clock errors.* The equations are simplified by leaving out the speed of light, or other conversion factor. Note: δ denotes an error variable.

$$\rho_k^i = R_k^i + \delta t_k + \delta t^i + I_k^i + T_k^i + M_k^i + C_k^i + v_k^i \tag{61}$$

ρ_k^i	Pseudorange of satellite i calculated by receiver k
R_k^i	True range between satellite i and antenna k
δt_k	Receiver k clock range error
δt^i	Satellite i clock range error
I_k^i	Ionospheric delay range between satellite i and antenna k
T_k^i	Tropospheric delay range between satellite i and antenna k
M_k^i	Multipath delay range between satellite i and antenna k
C_k^i	Code tracking loop range error of receiver k on the i^{th}satellite signal
v_k^i	White noise range error of receiver k on the i^{th} satellite signal

It is broken into the sum of the true range, several modelled errors, and a white noise term. The white noise term lumps all unmodeled errors, too small to model separately, into one random process, assumed white and Gaussian for modelling purposes. Normally, a pseudorange is corrected, open-loop, before it is used in any navigation algorithm. These corrections are described in particular for each error in later sections. This process can even be performed on the range itself so that every term is either an error or relatively small. The result is shown in Equation (62):

$$\delta\rho_k^i = \delta R_k^i + \delta t_k + \delta t^i + \delta I_k^i + \delta T_k^i + M_k^i + C_k^i + v_k^i \tag{62}$$

Simulation models for these error states have been developed for use at AFIT during the past decade [6, 9] and are described in the next sections.

B.3 Clock Bias and Drift

The largest error in the range calculation is due to clock differences. Mathematically, only the difference between the satellite clock and user clock determines the error magnitude, but for modelling purposes both the satellite and user clocks are compared to some arbitrary time standard, GPS time. The quartz crystal in the average user clock can generate large errors, especially if the receiver has been turned off for a long period of time. Two states are used to generate a representative user clock error. The first is an offset, or bias, from GPS time and the second is a drift rate. These errors are sometimes referred to as phase and frequency errors for obvious reasons. The initial offset and drift rate are modelled as random constants with initial conditions shown below [9]. In Kalman filtering it is often assumed that *all error variables are zero-mean Gaussian processes.*

$$\begin{bmatrix} \delta\dot{t}_k \\ \delta\ddot{t}_k \end{bmatrix} = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} \delta t_k \\ \delta\dot{t}_k \end{bmatrix} \tag{63}$$

$$\hat{\mathbf{x}}_{\delta t_k}(t_o) = \begin{bmatrix} 0 \text{ ft} \\ 0 \text{ ft/sec} \end{bmatrix} \tag{64a}$$

$$\hat{\mathbf{P}}_{\delta t_k}(t_o) = \begin{bmatrix} 9e14 \text{ ft}^2 & 0 \\ 0 & 9e10 \text{ ft}^2/\text{sec}^2 \end{bmatrix} \tag{64b}$$

Equation (63) is a stochastic differential equation and is of the form: $\dot{\mathbf{x}} = \mathbf{F}\mathbf{x} + \mathbf{B}\mathbf{u} + \mathbf{G}\mathbf{w}$ with $\mathbf{u} = \mathbf{0}$ and $\mathbf{w} = \mathbf{0}$. The e in Equation (64b) denotes a multiplication of the preceding number by ten raised to the power of the following number. All bold lower case letters are column vectors and upper case bold letters are matrices.

Each satellite clock is actually a combination of four atomic clocks, two cesium and two rubidium. This results in an extremely stable and accurate clock. The GPS control segment attempts to keep the entire constellation as synchronized as possible to GPS time. There are three correction values for each satellite in the navigation message. These values are updated regularly and correct for as much error as possible, but there is still some uncompensated error. The model for the

SV clock corresponds to the error after the clock corrections in the navigation message have been applied. The resulting SV clock error is small, usually accurate to one meter (the distance light travels in about three nanoseconds), and slowly varying, assumed constant during each experiment for this simulation. It is modelled as a random constant with initial conditions shown below [9]. It is assumed there are no correlations among SV clock errors.

$$\delta t^i = 0 \tag{65}$$

$$\widehat{x}_{\delta t^i}(t_o) = 0 \text{ ft} \tag{66a}$$

$$\hat{P}_{\delta t^i}(t_o) = 25 \text{ ft}^2 \tag{66b}$$

B.4 Ionospheric Delay

The ionosphere is a layer in the atmosphere above the stratosphere that contains an abundance of charged particles. These particles have an effect on electromagnetic (EM) waves that pass through it. In the case of code modulated on an EM wave, the charged particles delay the information. The delay is directly related to the total number of electrons in the path of the signal. This is named the total electron count (TEC). Another problem is that TEC is highly varying and dependent on factors such as time of day, latitude, and solar activity. Uncompensated, this delay can exceed 150 feet. An extreme amount of work [5, 15] has been accomplished in this area and there are time-of-day and seasonal models which can correct for some of this error. This is how the delay is reduced in most civilian receivers. The best solution is to use a two-frequency receiver. By performing a differencing operation, described in a later section, the ionospheric delay is calculated fairly well. There is some error left, up to ten feet, which can be modelled as a first-order, Gauss-Markov process [9]. Notice it is slowly varying with a time constant of 1500 seconds. Do not confuse the Dirac delta function notation in Equation (69) with an error. The properties of a Gauss-Markov process allow the noise statistics to be *completely described by the mean, covariance, and covariance kernel.*

In first-order processes this simplifies to the mean, variance, and correlation time. This property is exercised throughout this work.

$$\delta\dot{I}^i = \frac{-1}{1500}\delta I^i + w_{\delta I} \tag{67}$$

$$\widehat{x}_{\delta I}(t_o) = 0 \text{ ft} \tag{68a}$$

$$\hat{P}_{\delta I}(t_o) = 3 \text{ ft}^2 \tag{68b}$$

$$E[w_{\delta I}(t_p)w_{\delta I}(t_q)^T] \quad = \quad .004\ \delta(t_p - t_q)\ \text{ft}^2/\text{sec}^2 \tag{69}$$

B.5 Tropospheric Delay

The troposphere is the layer of atmosphere closest to the Earth. It is usually defined as the first eleven kilometers off the surface and contains over 90 percent of the mass of the entire atmosphere. This layer tends to bend and diffract EM waves passing through it. The effect is to cause another delay in the GPS signal. This delay can be upwards of 80 feet if uncompensated. The magnitude is a function of satellite elevation angle, altitude, and weather conditions. Open-loop compensation can reduce the error to a point where it can be modeled as a first-order, Gauss-Markov process [9]:

$$\delta\dot{T}^i = \frac{-1}{500}\delta T^i + w_{\delta T} \tag{70}$$

$$\widehat{x}_{\delta T}(t_o) = 0 \text{ ft} \tag{71a}$$

$$\hat{P}_{\delta T}(t_o) = 1 \text{ ft}^2 \tag{71b}$$

$$E[w_{\delta T}(t_p)w_{\delta T}(t_q)^T] \quad = \quad .004\ \delta(t_p - t_q)\ \text{ft}^2/\text{sec}^2 \tag{72}$$

B.6 Multipath Delay

When reflective surfaces are near either the satellite or user, the signal distortion known as multipath can occur. This effect is the superposition of one, or more, delayed instances of the same signal out of phase with each other. This superposition of out-of-phase signals causes the original signal to seem delayed. The best way to deal with multipath is to minimize its effect. It is important to note the GPS signal is right hand circularly polarized (RHCP). When a signal reflects off a surface with a higher dielectric constant than the medium through which it is currently travelling, it reverses polarization. Because air has an extremely low dielectric constant, the polarization is almost always reversed when a reflection occurs. GPS receiver antennas normally are designed to receive only RHCP signals. Signals reflected an odd number of times are attenuated when received at the antenna, in addition to the reflection loss incurred. The GPS signal structure is inherently resistant to the effect of multipath, but not immune. A reflected signal delayed by more than 1.5 chip widths[10] will be automatically rejected by the receiver, so signals with less delay than 1.5 chips are responsible for the error. Large errors can occur, especially in areas with large structures. Multipath has even been found to occur off the wings of aircraft or off the ground when landing [12]. Many antenna designs highly attenuate signals which originate low on the horizon. This reduces the effect of many reflected signals. Other antennas adaptively change their reception pattern to increase the strength of signals originating from the direction of known satellites. This type of antenna is called a Controlled Radiation Pattern Antenna (CRPA).

The average military aircraft does not have a CRPA and neither does a passenger aircraft; therefore, there is still a desire to model the effects of multipath. The multipath error has a time-varying correlation time, due to vehicle dynamics, changing reflective environment, and relative geometries. This is a very difficult error to simulate. A simple first-order lag with an unchanging time constant can be used to create a representative effect. This is not intended to recreate a real multipath error. The main focus of this addition is to input an unknown error that is not

[10]Chip width refers to the width of a bit of the pseudorandom code that is used to spread the GPS carrier signal. For C/A-code 1.5 chip widths corresponds to 1.47 μs, or 339 m, and for P-code, 147 ns, or 34 meters.

compensated by DGPS or other technique and see the effects on estimation. In fact, multipath effects are generally increased by differencing. The characteristics chose are shown below:

$$\dot{M}_k^i = -\frac{1}{10} M_k^i + w_M \tag{73}$$

$$\hat{x}_M(t_o) = 0 \text{ ft} \tag{74a}$$

$$\hat{P}_M(t_o) = 4 \text{ ft}^2 \tag{74b}$$

$$E[w_M(t_p) w_M(t_q)^T] \quad = \quad .2\ \delta(t_p - t_q)\ \text{ft}^2/\text{sec}^2 \tag{75}$$

Notice that this is a non-zero mean error. It was found that during the landing approach of a Boeing 747, a GPS receiver had a mean error of one meter due to multipath [12]. In the case of multipath, only positive values, delays, can occur in the real world. This property deserves special attention. Normally, the total value is estimated open loop and the error of that estimate is generated by the truth model. This way it can be justified that the error is Gaussian and zero-mean. In the case of the multipath state, this is not true. In simulation, this truth state *must* be forced to stay positive. Now the filter is simply estimating a variable which happens to always be positive. Since it is unaware of the ensemble statistics, it doesn't mind at all that, on a particular run, the state never goes negative. A better solution to this problem may be the addition of a bias state to shift the error positive. Also, a time-varying correlation time and a time-varying noise strength would produce a more realistic model [17]. Therefore, a very representative multipath model would have four states: bias, correlation time, noise strength, and multipath error.

B.7 Code Tracking Loop Error

All GPS receivers use tracking loops to maintain lock on the satellite transmissions. The dynamically-correlated error committed by a typical phase locked loop can be modelled as a first order lag driven by white noise. This technology continues to improve and become less expensive.

The numbers in the following equations may be slightly outdated, but they are kept to maintain a baseline of simulation performance here at AFIT [9]. Each channel of the receiver is independent of every other channel. This state is sometimes combined with the white noise term, because of its fast correlation time.

$$\dot{C}_k^i = -C_k^i + w_C \tag{76}$$

$$\widehat{x}_C(t_o) = 0 \text{ ft} \tag{77a}$$

$$\hat{P}_C(t_o) = 1 \text{ ft}^2 \tag{77b}$$

$$E[w_C(t_p)w_C(t_q)^T] \;=\; 2\,\delta(t_p - t_q)\ \text{ft}^2/\text{sec}^2 \tag{78}$$

B.8 Ephemeris Error

Satellite orbits are forever changing. They are constantly being perturbed by numerous phenomena, such as third body effects, magnetic anomalies, etc. The accuracy of the satellite positions is limited by the precision of the GPS control segment here on Earth. The ephemeris data is updated hourly and to a mean accuracy of 3-4 meters [4]. It should be noted that this figure is dated and the tracking filters are constantly improving. The discrete updates performed by the control segment are not easily modelled in a KF and usually not modelled in simulation. Orbit velocity errors are also tracked, but since we are not using range rate measurements, position error states are sufficient. Each satellite uses three states, one for each dimension. These states are unobservable when they are orthogonal to the SV line of sight (LOS) vector. Sometimes the ephemeris error is modelled as a single LOS error state. While this is more efficient, it may not be representative of the real world for long periods of time. This error is fundamentally different from all the other errors discussed in this section. Notice that is does not appear in the raw PR equation. That is

because the ephemeris error does not affect the time it takes the signal to travel or the calculation of that time. It is for the time of transit that the PR is actually calculated. The ephemeris error affects the solution when the locations of the satellites are taken to be known precisely, but are in fact incorrect. The sign of the error is defined in Equation (82):

$$\begin{bmatrix} \delta\dot{r}_x^i \\ \delta\dot{r}_y^i \\ \delta\dot{r}_z^i \end{bmatrix} = \begin{bmatrix} \frac{-1}{1500} & 0 & 0 \\ 0 & \frac{-1}{1500} & 0 \\ 0 & 0 & \frac{-1}{1500} \end{bmatrix} \begin{bmatrix} \delta r_x^i \\ \delta r_y^i \\ \delta r_z^i \end{bmatrix} + \mathbf{w}_{\delta r} \tag{79}$$

$$\hat{\mathbf{x}}_{\delta r}(t_o) = \mathbf{0} \text{ ft} \tag{80a}$$

$$\hat{\mathbf{P}}_{\delta r}(t_o) = \begin{bmatrix} 30 & 0 & 0 \\ 0 & 30 & 0 \\ 0 & 0 & 30 \end{bmatrix} \text{ ft}^2 \tag{80b}$$

$$E[\mathbf{w}_{\delta r}(t_p)\mathbf{w}_{\delta r}(t_q)^T] = \begin{bmatrix} .4 & 0 & 0 \\ 0 & .4 & 0 \\ 0 & 0 & .4 \end{bmatrix} \delta(t_p - t_q) \text{ ft}^2/\text{sec}^2 \tag{81}$$

$$\mathbf{r}^i = \hat{\mathbf{r}}^i + \boldsymbol{\delta}\mathbf{r}^i \tag{82}$$

$\mathbf{r}^i$ True position of satellite i
$\hat{\mathbf{r}}^i$ Position in ephemeris of satellite i

BIBLIOGRAPHY

[1] Athans, M. and C. B. Chang. *Adaptive Estimation and Parameter Identification Using a Multiple Model Estimation Algorithm.* Technical Note 1976-28, Lexington, MA: Lincoln Lab, MIT, June 1976. ESD-TR-184.

[2] Britting, Kenneth R. *Inertial Navigation Systems Analysis.* New York:Wiley-Interscience, 1971.

[3] Carlson, Neal A. and Jr. C. M. Neily. *Distributed Kalman Filter Architectures.* Final Technical Report, AFWAL-TR-87-1181, Air Force Avionics Laboratory, Wright-Patterson AFB, OH, June 1987.

[4] Crum, Jeffrey D. and Jr. Ronald T. Smetek. Welcome to the Machine: An Overview of GPS Master Control Station Anomaly Detection and Resolution Techniques. Navigation: Journal of The Institute of Navigation, Vol. 44, No. 2, Summer 1997.

[5] Daniell, Jr. R. E. "Improved GOS Ionospheric Correction Modeling," SBIR Topic AF95-105 (Phase I), Phillips Laboratory, Hanscom AFB, MA, December 1995.

[6] Gray, Robert A. *An Integrated GPS/INS/BARO and Radar Altimeter System for Aircraft Precision Approach Landings.* MS Thesis, AFIT/GE/ENG/94D-14, School of Engineering, Air Force Institute of Technology, Wright-Patterson AFB, OH, December 1994.

[7] Johnston, Kenneth. *Analysis of Radio Frequency Interference Effects on a Modern Coarse/Acquistion Code Global Positioning System Receiver.* MS Thesis, AFIT/GSO/ENG/99M-02, School of Engineering, Air Force Institute of Technology, Wright-Patterson AFB, OH, March 1999.

[8] Lewantowicz, Z. H. and D. W. Keen. "Graceful Degradation of GPS/INS Performance with Fewer Than Four Satellites," The Institute of Navigation, National Technical Meeting, Phoenix, AZ, 269-275, January 1991.

[9] Lewantowicz, Zdzislaw H. and Randall N. Paschall. "Deep Integration of GPS, INS, SAR, and Other Sensor Information", Unpublished technical article, 1995.

[10] Maybeck, Peter S. *Stochastic Models, Estimation, and Control, I.* New York: Academic Press, Inc., 1979. Republished, Arlington, VA: Navtech, 1994.

[11] Maybeck, Peter S. *Stochastic Models, Estimation, and Control, II.* New York: Academic Press, Inc., 1982. Republished, Arlington, VA: Navtech, 1994.

[12] Murphy, Tim, Robert Snow, and Michael Braash. "GPS Multipath on Large Commercial Air Transport Airframes. Navigation," Journal of The Institute of Navigation, Vol. 43, No. 4, Winter 1996- 1997.

[13] Musick, Stanton H. and Neil Carlson. *User's Manual for a Multimode Simulation for Optimal Filter Evaluation (MSOFE).* Air Force Avionics Laboratory, Wright-Patterson AFB, OH, April 1990. AFWAL-TR-88-1136.

[14] Papoulis, A., *Probability , Random Variables, and Stochastic Processes.* McGraw-Hill, New York, 1965.

[15] Parkinson, Bradford W. *Global Positioning System: Theory and Applications.* Washington DC: American Institute of Aeronautics and Astronautics, Inc., 1996.

[16] Raquet, John F. *Development of a Method for Kinematic GPS Carrier-Phase Ambiguity Resolution Using Multiple Reference Receivers.* PhD Dissertation, Department of Geomatics Engineering, University of Calgary, Calgary, Alberta, Canada, May 1998.

[17] Raquet, John F. and G. Lachapelle. "Long Distance Kinematic Carrier-Phase Ambiguity Resolution using a Simulated Reference Receiver Network," The Institute of Navigation, National Technical Meeting, Kansas City, MO, September 1997.

[18] Ritland, John T. Impact of Inertial System Quality on GPS-Inertial Performance in Jamming Environment. American Institute of Aeronautics and Astronautics, Inc. 1987. 1459-1467.

[19] White, Nathan A. *MMAE Detection of Interference/Jamming and Spoofing in a DGPS-Aided INS.* MS Thesis, AFIT/GE/ENG/96D-21, School of Engineering, Air Force Institute of Technology, Wright-Patterson AFB, OH, December 1996.

Vita

Barry Joseph Vanek was born 10 June 1974 in Schuyler, Nebraska. Upon graduating from Schuyler Central High School in 1992, he entered the Reserve Officer Training Program at Iowa State University. During his time in Ames he spent a year as the program manager of the High-Altitude Balloon Experiments in Technology (HABET) project. He received an Air Force commission and Bachelor of Science Degree in Electrical Engineering in May 1997. From August 1997 to March 1999, he attended the Air Force Institute of Technology (AFIT) where he received a Master of Science Degree in Electrical Engineering with an emphasis in navigation and control systems. Following AFIT, he was assigned to the National Air Intelligence Center, Wright-Patterson AFB, OH.